UNIVERSITY COLLEGE LONDON

Quasi-experimental Studies in Applied Microeconomics

Mirko Draca

A thesis submitted to the Department of Economics of University College London

for the degree of Doctor of Philosophy,

London, September 2013.

Declaration

I, Mirko Draca, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

I declare that my thesis consists of 77,853 words.

None of the work presented in this dissertation has been put forward for any previous degree.

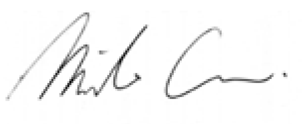

Mirko Draca

Acknowledgements

I thank my supervisor Steve Machin for his advice and support throughout.

I also thank my co-authors: Nick Bloom; Jordi Blanes i Vidal; Christian Fons-Rosen; John Van Reenen; and Robert Witt.

Abstract

In this dissertation I use a quasi-experimental approach across five different applied microeconomic studies. These studies are diverse in the range, covering police and crime, the political economy of lobbying, the effects of the minimum wage, and 'induced innovation' by firms in response to different incentives. However, each study outlines a comprehensive quasi-experimental approach that addresses potential threats to the given identification strategy. As a result, these studies provide credible, causal estimates of a number of important economic parameters including: the police-crime elasticity, the value of political connections among US Federal lobbyists, the impact of the minimum wage, and different incentives affecting technology adoption and innovation at the firm-level.

Impact Statement

The research in my dissertation has the capacity to make a number of impacts inside and outside academia. The primary impact outside academia relates to the first chapter which estimates the relationship between police and crime. The empirical evidence developed in this chapter can be used to inform policy decisions about increases or reductions in police resources. Another major impact outside of academia relates to chapter 2 on lobbying in Washington. The evidence in this chapter establishes the existence of a 'market for access' in the US lobbying industry and therefore provides a rationale for developing new regulatory frameworks for controlling money in US politics. Finally, the focus on Chinese import competition in chapter 4 has scientific importance within academia. It is one of the first studies to examine the impact of the early 2000s wave of Chinese imports on economic outcomes in developed economies.

TABLE OF CONTENTS

Contents

List of Tables

Chapter 5

List of Figures

OVERVIEW

In this dissertation I present five studies that use the 'quasi-experimental' or 'research design' approach to applied economic work. Specifically, each chapter addresses the issue of statistical causality and establishes a detailed identification strategy to validate the quasi-experiment being considered. In this overview I briefly summarise the main question being addressed in each study and the most important elements of the supporting identification strategy.

The work in chapter 1 ("Panic on the Streets of London") studies the link between police and crime, using the police deployment that followed the 2005 London terrorist attacks as the basis of a natural experiment. The reallocation of police after the attacks (principally, the 6-week-long, 35% increase in police presence across central London boroughs) is used to build a treatment-comparison group design that compares crime trends in different areas. Empirically, the ultimate result is an estimate of the police-crime elasticity of approximately -0.3. This estimate is in line with previous estimates but importantly is a more *precise* estimate than has been achieved before, in part due to the police deployment information used (which has not been used in previous, similar studies) and also due to the identification strategy.

The empirical strategy explicitly deals with the issue of correlated observable and unobservable shocks that may interfere with the interpretation of the police deployment effect. The intuition here is that post-attack changes in social and economic behavior not directly related to the police deployment may have affected crime independently. This chapter takes a number of steps to rule this possibility out. In particular, the empirical work shows that the drop in crime is closely timed with the shifts in police deployment – if the correlated shocks were driving crime trends then they would need to have a very unusual place-specific, 6-week duration to explain the police deployment effect.

Chapter 2 looks at the labour market for lobbyists in Washington and the role of political connections (defined in terms of past employment experience on the staff of Congressional politicians) as a determinant of revenues. This is a important policy question given recent concerns about the role of special interests in influencing political decisions. The empirical strategy studies what occurs to lobbyist revenues after an affiliated Senator or House Representative leaves the Congress. Practically, when an affiliated politician leaves the Congress then a lobbyist suffers a loss in 'politician-specific capital', that is, a premium that is received for their (the lobbyist's) expertise in the legislative preferences of an individual politician. The estimate provided in this chapter is that losing a connection to a Senator or senior House Representative leads to a 25% drop in revenues. As expected, this drop is even higher for connections to longer tenured politicians who sat on major committees prior to their exit from the Congress. Empirically, the work presented in this chapter addresses identification concerns about 'anticipated' exits by showing that the drop in revenues is highly discontinuous and concentrated around the exit point.

In the third chapter I address the impact of the minimum wage on firm profitability. A wave of research during the 1990s considered the impact of the minimum wage on employment and was not able to pin down the decisive negative effects that were widely expected in some quarters. This leaves open the question of how the wage effects of the policy are being transmitted through the economy? The empirical work in this chapter provides evidence that firm profitability may be one channel. Using two UK datasets – one on a large sample of companies and the other on the care homes industry – I find that wages increased and profits fell in firms after the introduction of the UK National Minimum Wage (NMW) in 1999. The magnitude of this effect is consistent with a 'no behavioral response' model whereby firms absorbed the increase in the wage bills with no (observable) countervailing adjustment.

Firms are also the focus of Chapter 4 which looks at the effects of rising Chinese import competition on technology adoption, productivity and innovation in a large sample of European firms. The major contribution of this chapter is that it shows that there is a *specific technology upgrading effect* of low-wage country import competition. Importantly, this upgrading effect is not just a reallocation effect that operates via a more intense selection effect with lower tech firms and products exiting the market and lifting average technology levels. Instead, this chapter identifies a distinctive *within-firm* technology effect that is equivalent in magnitude to the reallocation effect. This finding is validated using a quasi-experiment based on the entry of China into the World Trade Organisation (WTO). The accession of China into the WTO led to the relaxation of import quotas in the European textiles, clothing and footwear

(TCF) industries and therefore a large increase in import competition for a subset of manufacturing industries.

The final chapter follows up on this theme of ‘induced innovation’ by examining the effects of government demand for high-tech goods in the US, as represented by defense spending. This demand channel for induced innovation has not been considered previously and I show that the effect is large both in terms of a parameterized demand effect and in terms of the historical magnitude. Specifically, I use shifts in spending associated principally with the ‘Reagan build-up’ of the 1980s to track firm-specific shifts in demand. I am able to track these shifts by relating how the *product composition* of defense procurement spending impacted on individual firms. The demand effect of defense procurement is estimated as an elasticity of approximately 0.07, which is nearly twice the benchmark established for civilian-sourced demand. Finally, in terms of magnitudes, the contribution of defense procurement to innovation peaked during the early Reagan build-up, accounting for 11.4% of the total change in patenting intensity and 6.5% for R&D. This compares to a defense sector share in output of around 4%. The later defense cutbacks under Bush Senior and Clinton then curbed the growth in technological intensity by around 2%.

CHAPTER 1: PANIC ON THE STREETS OF LONDON: POLICE, CRIME AND THE JULY 2005 TERROR ATTACKS

Abstract

In this paper, we study the causal impact of police on crime, looking at what happened to crime and police before and after the terror attacks that hit central London in July 2005. The attacks resulted in a large redeployment of police officers to central London as compared to outer London – in fact, police deployment in central London increased by over 30 percent in the six weeks following the July 7 bombings, before sharply falling back to pre-attack levels. During this time, crime fell significantly in central relative to outer London. Study of the timing of the crime reductions and their magnitude, the subsequent sharp return back to pre-attack crime levels, the types of crime that were more likely to be affected and a series of robustness tests looking at possible biases all make us confident that our research approach identifies a causal impact of police on crime. The instrumental variable approach we use uncovers an elasticity of crime with respect to police of approximately -0.3 to -0.4, so that a 10 percent increase in police activity reduces crime by around 3 to 4 percent.

JEL Classifications*:* H00, H5, K42.

Keywords: Crime; Police; Terror attacks.

1.1 Introduction

Terrorism is arguably the single most significant topic of political discussion of the past decade. In response, a small economic literature has begun to investigate the causes and impacts of terrorism (see Krueger, 2007, for a summary or Krueger and Jitka Maleckova, 2003, for some empirical work). Terror attacks, or the threat thereof, have also been considered in research on one important area of public policy, namely the connections between crime and policing. Some recent studies (such as Di Tella and Schargrodsky, 2004 and Klick and Taborrak, 2005) have used terrorism-related events to look at the crime-police relationship since terror attacks can induce an increased police presence in particular locations. This deployment of additional police can be used, under certain conditions, to test whether or not increased police reduce crime.

In this paper, we also consider the crime-police relationship before and after a terror attack, but in a very different context to other studies by looking at the increased security presence following the terrorist bombs that hit central London in July 2005. Our application is a more general one than the other studies in that it covers a large metropolitan area following one of the most significant and widely known terror attacks of recent years. The scale of the security response in London after these attacks provides a potentially useful setting to examine the relationship between crime and police.

Moreover, and unlike the other studies in this area, we have very good data on police deployment. We can use these to identify the magnitude of the causal impact of police on crime.[1] A major strength of this paper is therefore that we are able to offer explicit instrumental variable-based estimates of the crime-police elasticity, which can be compared to other estimates like Levitt (1997), Corman and Mocan (2000) and Di Tella and Schargrodsky's (2004) implied

---

[1] Neither Di Tella and Schargrodsky (2004) nor Klick and Tabarrok (2005) have data on police activity.

elasticity.[2] In fact, the sharp discontinuity in police deployment that we able to identify using this data means we are able to pin down this causal relation between crime and police very precisely. The natural experiment that we consider also has some important external validity in the sense that it involves the deployment of a clear "deterrence technology" (that is, more police on the streets) rather than a measure of increased expenditures on police (e.g. as in Evans and Owens, 2007, or Machin and Marie (2011)). Arguably, this type of visible increase in police deployment is the main type of policy mechanism under discussion in public debates about the funding and use of police resources.

Furthermore, the effectiveness of police is important in the context of a large criminological literature that has generally failed to find significant impacts of police on crime, even in quasi-experimental studies. Sherman and Weisburd (1995) review some of the conclusions from this work. Gottfredson and Hirschi (1990: 270) state "no evidence exists that augmentation of police forces or equipment, differential police strategies, or differential intensities of surveillance have an effect on crime rates". Similarly emphatic arguments are made in Klockars (1983).

The focus of the current paper is on what happened to criminal activity following a large and unanticipated increase in police presence. The scale of the change in police deployment that we study is much larger than in any of the other work in the crime-police research field. Indeed, results reported below show that police activity in central London increased by over 30 percent in the six weeks following the July 7 bombings as part of a police deployment policy stylishly titled "Operation Theseus" by the authorities. This police intervention represented the deployment of a very strong deterrence technology. The coverage of police was more sustained, widespread and complete than in the previous studies in the existing literature. We therefore view the scale of this change as important in addressing the paradox of the criminology literature

---

[2] Whilst the Levitt (1997) paper is well known and widely cited, Justin McCrary's (2002) comment highlights some concerns about the data and the approach used (see also Levitt's, 2002, response).

discussed above where it proves hard to detect crime reductions linked to increased police presence. This is particularly the case since, during the time period when police presence was heightened, crime fell significantly in central London relative to outer London. Both the timing of the crime reductions and the types of crime that were more affected make us confident that this research approach identifies a causal impact of police on crime. Moreover, when police deployments returned to their pre-attack levels some six weeks later, the crime rate rapidly returned to its pre-attack level. Exploiting these sharp discontinuities in police deployment, we estimate an elasticity of crime with respect to police of approximately -0.3 to -0.4, so that a 10 percent increase in police activity reduces crime by around 3 to 4 percent. Furthermore, we are unable to find evidence of either temporal or contemporaneous spatial displacement effects arising from the six-week police intervention.

A crucial part of identifying a causal impact in this type of setting is establishing the exclusion restriction that terrorist attacks affect crime through the post-attack increase in police deployment, and not via other observable and unobservable factors correlated with the attack or shock. Establishing this is important to generate credibility that our findings inform the crime-police debate rather than being just about an episode where a terror attack occurred. In this regard, the police deployment data we use are invaluable as their availability make it possible to distinguish the impact of police on crime from any general impact of the terrorist attack. In particular, the research design features two discontinuities related to the police intervention. The first is the introduction of the geographically focused police deployment policy in the week of the terrorist attack. The immediate period surrounding the introduction of the policy was also characterized by a series of potentially correlated observable and unobservable shocks related to the attack. In contrast, the second discontinuity associated with the withdrawal of the policy occurred in a very different context. In this case, the observable and unobservable shocks associated with the attack were still in effect and dissipating gradually. Crucially though, the police deployment was discretely "switched off" after a six week period and we observe an increase in crime that is exactly timed with this change. Thus, we argue that is difficult to attribute such a clear change in crime rates to observable and unobservable shocks arising from the terrorist attacks. If these types of shocks significantly affected crime rates, we would expect this to continue even as the police deployment was withdrawn. Indeed, an interesting feature of our empirical results is just how clearly and definitively crime seems to respond to a police presence.

The rest of the paper is organized as follows. Section I describes the events of July 2005 and goes over the main modeling and identification issues. In Section II we describe the data and provide an initial descriptive analysis. Section III presents the statistical results and a range of additional empirical tests. Section IV concludes.

1.2 Crime, Police and the London Terror Attacks

1.21 The Terror Attacks

In July 2005 London's public transport system was subject to two waves of terror attacks. The first occurred on Thursday 7th July and involved the detonation of four bombs. The 32 boroughs of London are shown in Figure 1. Three of the bombs were detonated on London Underground (the tube) train carriages near the stations of Russell Square (in the borough of Camden), Liverpool Street (Tower Hamlets) and Edgware Road (Kensington and Chelsea). A fourth bomb was detonated on a bus in Tavistock Square, Bloomsbury (Camden). The second wave of attacks occurred two weeks later on the 21st July, consisting of four unsuccessful attempts at detonating bombs on trains near the underground stations of Shepherds Bush (Kensington and Chelsea), the Oval (Lambeth), Warren Street (Westminster) and on a bus in Bethnal Green (Tower Hamlets). Despite the failure of the bombs to explode, this second wave of attacks caused much turmoil in London. There was a large manhunt to find the four men who escaped after the unsuccessful attacks and all of them were captured by 29th July.

1.22 Terror Attacks, Crime and Correlated Shocks

Di Tella and Schargrodsky (2004) were first to use police allocation policies in the wake of terror attacks to circumvent the endogeneity problem of crime and police. Using a July 1994 terrorist attack that targeted the main Jewish center in Buenos Aires, they show that motor vehicle thefts fell significantly in areas where extra police were subsequently deployed compared to areas several blocks away that did not receive extra protection. Their effect is large (approximately a 75 percent reduction in thefts relative to the comparison group) but also extremely local with no evidence that the police presence reduced crime one or two blocks away from the protected areas. Another study by Klick and Tabarrok (2005) uses terror alert levels in Washington DC to make inferences about the crime-police relationship. The deployments they consider cover a more general area but (as already discussed) in the end are speculative since they are not able to quantify them with data on police numbers or hours.

Both of these papers touch on the issue of correlated shocks to observables and unobservables. However, in our case of London this could be a greater concern since the terrorist attacks were a more significant, dislocating event for the city. Therefore, in thinking about the question of correlated shocks, it is helpful to first consider a basic equation, specified in levels, that describes the determinants of the crime rate in a set of geographical areas (in our case, London boroughs) over time:$= \alpha + \delta P_{jt} +$

$$C_{jt} = \alpha + \delta P_{jt} + \lambda X_{jt} + \mu_j + \tau_t + \varepsilon_{jt} \qquad (1.1)$$

where C_{jt} is the crime rate for borough j in time period t, P_{jt} the level of police deployed and X_{jt} is a vector of control variables that could be comprised of observable or unobservable elements. The next set of terms are: μ_j, a borough level fixed effect; τ_t, a common time effect (for example, to capture common weather or economic shocks) and ε_{jt} is a random error term. In

(1) t denotes weeks as we estimate weekly crime equations, in which we are careful to recognise that crime displays a strong seasonal persistence.[3]

Consider a seasonally differenced version of equation (1), where the dependent variable is the change in the area crime rate relative to the rate in the same week of the previous year. This is highly important in crime modeling since crime is strongly persistent across areas over time. In practical terms, differencing eliminates the borough-level fixed effect yielding:

$$(C_{jt} - C_{j(t-k)}) = \delta(P_{jt} - P_{j(t-k)}) + \lambda(X_{jt} - X_{j(t-k)}) + (\tau_t - \tau_{t-k}) + (\varepsilon_{jt} - \varepsilon_{j(t-k)}) \qquad (1.2)$$

note that the $\tau_t - \tau_{t-k}$ difference term can now be interpreted as the year-on-year change in factors that are common across all of the areas. By expressing this equation more concisely we can make the correlated shocks issue explicit as follows:

$$\Delta_k C_{jt} = \delta\Delta_k P_{jt} + \lambda\Delta_k X_{jt} + \Delta_k \tau_t + \Delta_k \varepsilon_{jt} \qquad (1.3)$$

where Δ is a difference operator with k indexing the order of the seasonal differencing.

Using this framework we can carefully consider how a terrorist attack – which we can denote generally as Z - affects the determinants of crime across areas. Following the argument in the papers discussed above, the terror attack Z affects ΔP_{jt}, shifting police resources in a way that one can hypothesise is unrelated to crime levels. This hypothesis is, of course, a crucial aspect of identification that needs serious consideration. For example, it is possible that Z could affect the elements of ΔX_{jt} creating additional channels via which terrorist attacks could influence crime rates.

What are these potential impacts or channels? The economics of terrorism literature stresses that the impacts of terrorism can be strong, but generally turn out to be temporary (OECD, 2002; Bloom, 2009) in that economic activity tends to recover and normalize

[3] These types of effects could prevail where seasonal patterns affect different boroughs with varying levels of intensity. For example, the central London boroughs are more exposed to fluctuations due to tourism activity and exhibit sharper seasonal patterns with respect to crime.

itself fairly rapidly. Of course, a sharp but temporary shock would still have ample scope to intervene in our identification strategy by affecting crime in a way that is correlated with the police response. In particular, three channels demand consideration. First there is the physical dislocation caused by the attack. A number of tube stations were closed and many Londoners changed their mode of transport after the attacks (e.g. from the tube to buses or bicycles). This would have reshaped travel patterns and could have affected the potential supply of victims for criminals in some areas. Secondly, the volume of overall economic activity was affected. Studies on the aftermath of the attack indicate that both international and domestic tourism fell after the attacks, as measured by hotel vacancy rates, visitor spending data and counts of domestic day trips (Greater London Authority, 2005). Finally, there may be a psychological impact on individuals in terms of their attitudes towards risk. As Becker and Rubinstein (2004) outline, this influences observable travel decisions as well as more subtle unobservable behavior.

To summarize, we think of these effects as being manifested in three elements of the X_{jt} vector outlined above:

$$X_{jt} = [X_{jt}^1, X_{jt}^2, \theta_{jt}] \qquad (1.4)$$

In (4), X_{jt}^1 is a set of exogenous control variables (observable to researchers), that includes observable factors such as area-level labour market conditions that change slowly and are unlikely to be immediately affected by terrorist attacks (if at all). The second X_{jt}^2 vector represents the observable factors that change more quickly and are therefore vulnerable to the dislocation caused by terrorist attacks. As discussed above, here we are thinking primarily of factors such as travel patterns that could influence the potential supply of victims to crime across areas. The final element θ_{jt} then captures an analogous set of unobservable factors that are susceptible to change due to the terrorist attack. In the spirit of Becker and Rubinstein's (2004) discussion, the main factor to consider here is fear or how individuals handle the risks associated with terrorism. For example, it is plausible that, in the wake of the attacks, commuters in London became more vigilant to suspicious activity in the transport system and in public spaces. This vigilance would have been focused mainly on potential terrorist activity, but one might expect that this type of cautious behaviour could have a spillover onto crime.

The implications of these correlated shocks for our identification strategy can now be clearly delineated. For our exclusion restriction to hold it needs to be shown that the terrorist attack Z affected the police deployment in a way that can be separately identified from Z's effect on other observable and unobservable factors that can influence crime rates. Practically, we show this later in the paper by mapping the timing and location of the police deployment shock and comparing it to the profiles of the competing observable and unobservable shocks.

1.23 Possible Displacement Effects

Another issue that could affect identification is crime displacement. Since the police intervention affected the costs of crime across locations and time, it may be that criminals take these changes into account and adjust their behavior. This raises the possibility that criminal activity was either diverted into other areas (e.g. the comparison group of boroughs) during Operation Theseus or postponed until after the extra police presence was withdrawn. The implication is that simple differences-in-differences estimates of the police effect on crime would be upwardly biased if these offsetting spatial displacement effects were not taken into account. Temporal displacement can have the opposite effect and we discuss this more in the final empirical section.

1.3 Data Description and Initial Descriptive Analysis

1.31 Data

We use daily police reports of crime from the London Metropolitan Police Service (LMPS) before and after the July 2005 terrorist attacks. Our crime data cover the period from 1st January 2004 to 31st December 2005 and are aggregated up from ward to borough level and from days to weeks over the two-year period. There are 32 London boroughs as shown on the map in Figure 1.1.[4] There are also monthly borough level data available over a longer time period that we use for some robustness checks.

The basic street-level policing of London is carried out by 33 Borough Operational Command Units (BOCUs), which operate to the same boundaries as the 32 London borough councils apart from one BOCU which is dedicated to Heathrow Airport. We have been able to put together a weekly panel covering 32 London boroughs over two years giving 3,328 observations. Crime rates are calculated on the basis of population estimates at borough level, supplied by the Office of National Statistics (ONS) online database.[5]

The police deployment data are at borough level and were produced under special confidential data-sharing agreements with the LMPS. The main data source used is CARM (Computer Aided Resource Management), the police service's human resource management system. This records hours worked by individual officers on a daily basis. We aggregate the deployment data to borough-level since the CARM data is mainly defined at this level. However, there is also useful information on the allocation of hours worked by incident and/or police operation.[6] While hours worked are available according to officer rank, our main hours measure is based on total hours worked by all officers in the borough adjusted for this reallocation effect. In addition to crime and deployment, we have also obtained weekly data on tube journeys for all stations from Transport for London (TFL). It is daily borough-level data aggregated up to weeks based on entries into and exits from tube stations. Finally, we also use data from the UK Labour Force Survey (LFS) to provide information on local labour market conditions.

[4] The City of London has its own police force and so this small area is excluded from our analysis.

[5] Web Appendix Table A1 shows some summary statistics on the crime data.

[6] Since the CARM information is also used for calculating police pay it is considered a very reliable measure of police activity. We gained access to this data after repeated inquiries to the MPS. The main condition for access was that we not reveal any strategic information about ongoing or individual, borough-specific police deployment policies.

1.32 Initial Approach

Our analysis begins by looking at what happened to police deployment and crime before and after the July 2005 terror attacks in London using a differences-in-differences approach. This rests upon defining a treatment group of boroughs in central and inner London where the extra police deployment occurred and comparing their crime outcomes to the other, non-treated boroughs. The police hours data we use facilitates the development of this approach, with two features standing out. First, the data allow us to measure the increase in total hours worked in the period after the attacks. The increase in total hours was accomplished through the increased use of overtime shifts across the police service and this policy lasted approximately six weeks. Secondly, the police data contain a special resource allocation code denoted as Central Aid. This code allows us to identify how police hours worked were geographically reallocated over the six-week period. For example, we can identify how hours worked by officers stationed in the outer London boroughs were reallocated to public security duties in central and inner London. The extra hours were mainly reallocated to the boroughs of Westminster, Camden, Islington, Kensington and Chelsea, and Tower Hamlets, with individual borough allocations being proportional to the number of tube stations in the borough.[7] These boroughs either contained the sites of the attacks or featured many potential terrorist targets such as transport nodes or significant public spaces. Using these two features of the data we are able to define a treatment group comprised of the five named boroughs. A map showing the treatment group is given in Figure 1.1. In most of the descriptive statistics and modeling below we use all other boroughs as the comparison group in order to simplify the analysis.

What did the extra police deployment in the treated boroughs entail? The number of mobile police patrols were greatly increased and officers were posted to guard major public

---

[7] We say "mainly reallocated" due to the fact that some mobile patrols crossed into adjacent boroughs and because some bordering areas of boroughs were the site of some small deployments. A good case here is the southern tip of Hackney borough (between Islington and Tower Hamlets). However, the majority of Hackney was not treated by the policy (since this borough is notoriously lacking in Tube station links) so we exclude it from the treatment group.

spaces and transport nodes, particularly tube stations. In areas of central London where many stations were located this resulted in a highly visible police presence and this is confirmed by public surveys conducted at the time.[8] This high visibility potentially exerted a deterrent effect on public, street-level crimes such as thefts and violent assault. We test for this prediction in the empirical work.

1.33 Basic Differences-in-Differences

In Table 1.1 we compare what happened to police deployment and to total crime rates before and after the July 2005 terror attacks in the treatment group boroughs as compared to all other boroughs. Police deployment is measured in a similar way to crime rates, that is, we normalize police hours worked by the borough population. Following the discussion in Section 2 we define the before and after periods in year-on-year, seasonally adjusted terms. This ensures that we are comparing like-with-like in terms of the seasonal effects prevailing at a given time of the year. For example, looking at Table 1.1 the crime rate of 4.03 per 1000 population in panel B represents the treatment group crime rate in the period from the 8th of July 2004 until the 19th of August 2004. The post-period or "policy on" period then runs from July 7th 2005 until August 18th 2005 with a crime rate of 3.59.[9] Thus by taking the difference between these "pre" and "post" crime rates we are able to derive the year-on-year, seasonally adjusted change in crime rates and police hours. These are then differenced across the treatment (T = 1) and comparison (T = 0) groups to get the customary differences-in-differences (DiD) estimate.

The first panel of Table 1.1 shows unconditional DiD estimates for police hours. It is clear that the treatment boroughs experienced a very large relative change in police deployment. Per capita hours worked increased by 34.6 percent in the DiD (final row, column 3). Arguably, the *composition* of this relative change is almost as important for our experiment as

---

[8] Table 1.A2 of the chapter Appendix reports the results of a survey of London residents in the aftermath of the attacks. Approximately 70 percent of respondents from inner London attested to a higher police presence in the period since the attacks. The lower percentage reported by outer London residents also supports the hypothesis of differential deployment across areas.

[9] The one day difference in calendar date across years ensures we compare the same days of the week.

the scale. The relative change was driven by an increase in the treatment group (of 72.8 hours per capita) with little change in hours worked for the comparison group (only 2.2 hours more per capita). This was feasible because of the large number of overtime shifts worked. In practice, this means that, while there was a diversion of police resources from the comparison boroughs to the treatment boroughs, the former areas were able to keep their levels of police hours constant. Obviously, this *ceteris paribus* feature greatly simplifies our later analysis of displacement effects, since we do not have to deal with the implications of a zero-sum shift of resources across areas. The next panel of Table 1.1 deals with the crime rates. It shows that crime rates fell by 11.1 percent in the DiD (final row, column 6). Again, this change is driven by a fall in treatment group crime rates and a steady crime rate in the comparison group. This is encouraging since it is what would be expected from the type of shift we have just seen in police deployment.

Weekly police deployment and crime rates are shown in Figures 1.2a and 1.2b. Here we do two things. First, we normalize crime rates and police hours across the treatment and comparison groups by their level in week one of our sample (i.e. January 2004). This re-scales the levels in both groups so that we can compare their evolution over time. Secondly, we mark out the attack ("policy-on") period in 2005, along with the comparison period in the previous year. As Figure 1.2a shows, this reveals a clear, sharp discontinuity in police deployment. Police hours in the treatment group rise immediately after the attack and fall sharply at the end of the six weeks of Operation Theseus.

The visual evidence for the crime rate in Figure 1.2b is less decisive because the weekly crime rates are clearly more volatile than the police hours data. This is to be expected insofar as police hours are largely determined centrally by policy-makers, while crime rates are essentially the outcomes of decentralized activity. This volatility does raise the possibility that the fall in crime rates seen in the Table 1.1 DiD estimates may simply be due to naturally occurring, short-run time series volatility rather than the result of a policy intervention – a classic problem in the literature (Donohue, 1998). After the correlated shocks issue this is probably the biggest modeling issue in the paper and we deal with it extensively in the next section.

1.4 Statistical Models of Crime and Police

In this section we present our statistical estimates. We begin with a basic set of estimates and then move on to focus on specific issues to do with different crime types, timing, correlated shocks and displacement effects.

1.41 Statistical Approach

The starting point for the statistical work is a DiD model of crime determination. We have borough level weekly data for the two calendar years 2004 and 2005. The terror attack variable (Z as discussed above) is specified as an interaction term T_b*POST_t, where T denotes the treatment boroughs and POST is a dummy variable equal to one in the post-attack period.

In this setting the basic reduced form seasonally differenced weekly models for police deployment and crime (with lower case letters denoting logs) are:

$$\Delta_{52}p_{bt} = \alpha_1 + \beta_1 POST_t + \delta_1(T_b*POST_t) + \lambda_1\Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{1bt} \tag{1.5}$$

$$\Delta_{52}c_{bt} = \alpha_2 + \beta_2 POST_t + \delta_2(T_b*POST_t) + \lambda_2\Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{2bt} \tag{1.6}$$

Due to the highly seasonal nature of crime noted above, the equations are differenced across weeks of the year (hence the k = 52 subscript in the Δ_k differences). The key parameters of interest are the δ's, the seasonally adjusted differences-in-differences estimates of the impact of the terror attacks on police deployment and crime.

These reduced form equations can be combined to form a structural model relating crime to police deployment, from which we can identify the causal impact of police on crime. The structural equation is:

$$\Delta_{52}c_{bt} = \alpha_3 + \beta_3 POST_t + \delta_3\Delta_{52}p_{bt} + \lambda_3\Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{3bt} \tag{1.7}$$

where the variation in police deployment induced by the terror attacks identifies the causal impact of police on crime. The first stage regression is equation (1.5) above and so equation (1.7) is estimated by instrumental variables (IV) where the T_b*POST_t variable is used as the instrument for the change in police deployment. Here the structural parameter of interest, δ_3 (the coefficient on police deployment), is equal to the ratio of the two reduced form coefficients, so that $\delta_3 = \delta_2/\delta_1$.

Finally, note that in some of the reduced form specifications that we consider below we split the $POST_t*T_b$ into two distinct post 7/7 time periods so as to distinguish the "post-policy" period after the end of Operation Theseus. This term is added in order to directly to test for any persistent effect of the police deployment, and importantly to explicitly focus upon the second 'experiment' when police levels fell sharply back to their pre-attack levels. Thus the reduced forms in (1.5) and (1.6) now become:

$$\Delta_{52}p_{bt} = \alpha_4 + \beta_4 POST_t + \delta_{41}(T_b*POST_t^1) + \delta_{42}(T_b*POST_t^2) + \lambda_4\Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{4bt} \quad (1.8)$$

$$\Delta_{52}c_{bt} = \alpha_5 + \beta_5 POST_t + \delta_{51}(T_b*POST_t^1) + \delta_{52}(T_b*POST_t^2) + \lambda_5\Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{5bt} \quad (1.9)$$

In these specifications $POST_t^1$ represents the six-week policy period immediately after the July 7th attack when the police deployment was in operation while $POST_t^2$ covers the time period subsequent to the deployment until the end of the year (that is, from the 19th of August 2005 until December 31st 2005).[10] Also note that a test of $\delta_{41} = \delta_{42}$ (in the police equation, (8)) or $\delta_{51} = \delta_{52}$ (in the crime equation, (9)) amounts to a test of temporal variations in the initial six week period directly after July 7th as compared to the remainder of the year.

1.42 Basic Differences-in-Differences Estimates

Table 1.2 provides the basic reduced form OLS and structural IV results for the models outlined in equations (5)-(9). For comparative purposes, we specify three terms to uncover the differences-in-differences estimate. Specifically, in columns (1) and (5) we include an interaction term that uses the full period from July 7th 2005 to December 31st 2005 to measure the post-attack period (in the Table denoting T_b*POST_t from equations (5) and (6) as T*Post-Attack). The adjacent columns ((2)-(4) and (6)-(8)) then split this period in two with one interaction term for

[10] As we discuss later police deployment levels in London boroughs were returned to their pre-attack baselines after the end of Operation Theseus.

the six-week Operation Theseus period (denoting $T_b*POST_t^1$ from equations (8) and (9) as T*Post-Attack1) and another for the remaining part of the year (denoting $T_b*POST_t^2$ as T*Post-Attack2). As already noted, the second term is useful for testing whether there were any persistent effects of the police deployment or any longer-term trends in the treatment group after police deployment fell back to its pre-attack levels.

The findings from the unconditional DiD estimates reported earlier are confirmed in the basic models in Table 1.2. The estimated coefficient on T*Post-Attack1 in the reduced form police equation shows a 34.1 percent increase in police deployment during Operation Theseus, and there is no evidence that this persists for the rest of the year (i.e. the T*Post-Attack2 coefficient is statistically indistinguishable from zero). For the crime rate reduced form there is an 11.1 percent fall during the six-week policy-on period with minimal evidence of either persistence or a treatment group trend in the estimates for the T*Post-Attack2 variable.[11] Despite this we include a full set of 32 borough-specific trends in the specifications in columns (7) and (8) to test robustness. The crime rate coefficient for the Operation Theseus period is halved, but the interaction term is still significant indicating that there was a fall in crime during this period that was over and above that of any combination of trends.

The coincident nature of the respective timings of the increase in police deployment and the fall in crime suggests the increased security presence lowered crime. The final three columns of the Table therefore show estimates of the causal impact of increased deployment on crime. Column (11) shows the basic IV estimate where the post-attack effects are constrained to be time invariant. Columns (12) and (13) allow for time variation to identify a more local causal impact. The instrumental variable estimates are precisely determined owing to the strength of the first stage regressions in the earlier columns of the Table. The preferred estimate with time-varying

[11] Whilst we have seasonally differenced the data one may have concerns about possible contamination from further serial correlation. We follow Marianne Bertrand et al (2004) and collapse the data before and after the attacks and obtain extremely similar results: the estimate (standard error) based on collapsed data comparable to the T*Post-Attack 1 estimate in column (6) of Table 1.2 was -.112 (.027).

terror attack effects (reported in column (12)) shows an elasticity of crime with respect to police of around -0.32. This implies that a 10 percent increase in police activity reduces crime by around 3.2 percent. The magnitudes of these causal estimates are similar to the small number of causal estimates found in the literature (they are also estimated much more precisely in statistical terms because of the very sharp discontinuity in police deployment that occurred). Levitt's (1997) study found elasticities in the -0.43 to -0.50 range, while Corman and Mocan (2000) estimated an average elasticity of -0.45 across different types of offences and Di Tella and Schargrodsky (2004) reported an elasticity of motor vehicle thefts with respect to police of -0.33.

OLS estimates are reported in columns (9) and (10). The column labelled 'levels' shows a pooled cross-sectional regression with a strongly significant positive coefficient on the police deployment variable. In column (10) we show estimates from a seasonally-differenced version of this OLS regression, obtaining a negligible, insignificant coefficient. This reflects the limited year-on-year change in police hours to be found when the seasonal difference is taken.

1.43 Different Crime Types

So far the results use a measure of total crimes. However, heterogeneity of the overall effect by the type of crime is potentially important. The pattern of the impact by crime type is an important falsification exercise. The main feature of Operation Theseus was a highly visible public deployment of police officers in the form of foot and mobile patrols, particularly around major transport hubs. We could therefore expect any police effect to be operating mainly through a deterrence technology based on greater visibility, generating an increase in the probability of detection for crimes committed in or around public places. As a result, the crime effect documented in Tables 1.1 and 1.2 should be concentrated in crime types more susceptible to this type of technology.

We therefore estimated reduced form treatment effects across the six major crime categories defined by the Metropolitan Police – thefts, violent crimes, sexual offences, robbery, burglary and criminal damage. Separate estimates by crime type are reported in Table 1.3, which shows there to be important differences across these groups. The estimates show strongly significant effects for thefts and violent crimes. These are comprised of crimes such as street-level thefts (picking pockets, snatches, thefts from stores, motor vehicle-related theft and tampering) as well as street-level violence (common assault, harassment, aggravated bodily harm). Of considerable note is the lack of any effect for burglary. As a crime that mainly occurs at night and in private dwellings, this is arguably the crime category that is least susceptible to a public deterrence technology.

In Table 1.4 we aggregate these major categories into a group of crimes potentially susceptible to Operation Theseus (thefts, violent crimes and robberies) and a group of remaining non-susceptible crimes (burglary, criminal damage and sexual offences). The point estimate for our preferred susceptible crimes estimate is -0.131 (column 3, panel (I)) which compares to an estimate of -0.109 for total crimes in column (7) of Table 1.2. There is a much smaller (in absolute terms), statistically insignificant estimate of -0.033 for non-susceptible crimes (reported in column (3), panel (II), of Table 4). We therefore use this susceptible crimes classification as the main outcome variable in the remainder of our analysis. The estimated elasticity of susceptible crimes with respect to police deployment in the column (8) model of the Table is -0.38 and is again very precisely determined.

1.44 Timing

The previous section cited the volatility of the crime rates and timing in general as an important issue. Given that we are using weekly data, there is a need to investigate to what extent short-term variations could be driving the results for our policy intervention. To test this we take the extreme approach of testing every week for hypothetical or “placebo” policy effects. Specifically, we estimate the reduced form models outlined in equations (5) and (6) defining a single week-treatment group interaction term for each of the 52 weeks in our data. We then run 52 regressions each featuring a different week*T_b interaction and plot the estimated coefficient and confidence intervals. The major advantage of this is that it extracts all the variation and volatility from the data in a way that reveals the implications for our main DiD estimates. Practically, this exercise is therefore able to test whether our 6-week Operation Theseus effect is merely a product of time-series volatility or variation that is equally likely to occur in other sub-periods.

We plot the coefficients and confidence intervals for all 52 weeks in Figures 1.3a and 1.3b. Figure 1.3a shows the results for police hours repeating the clear pattern seen in Figure 1.2a of the police deployment policy being switched on and off. (Note that precisely estimated treatment effects in this graph are characterized by confidence intervals that do not overlap the zero line). The analogous result for the susceptible crime rate is then shown in Figure 1.3b. The decreases in crime are less dramatic than the increases in police hours, but the two clearly closely coincide in timing. It is interesting to note that the pattern of six consecutive weeks of significant, negative treatment effects in the crime rate is not repeated in any other period of the data *except* Operation Theseus. This is impressive as it shows that the effect of the policy intervention can be seen despite the noise and volatility of the weekly data.[12]

1.45 Correlated Shocks

The discussion of timing has a direct bearing on the issue of correlated shocks outlined in Section 1.2. In particular, it is important to examine the extent to which any shifts in correlated observables do or do not coincide in timing with the fall - and subsequent bounce back - in crime. The major observable variable we consider here concerns transport decisions and we study this

---

[12] As a further check on the issue of volatility we made use of monthly, borough-level crime data available from 2001 onwards (as the daily crime data we use to construct our weekly panel is only available since the beginning of 2004). These data allow us to examine whether there is a regular pattern of negative effects in the middle part of the year. In this exercise, we define year-on-year differences in susceptible crime for the July-August period over the a range of intervals: 2001-2002, 2002-2003, 2003-2004, and 2004-2005. The results are shown in chapter Appendix Table 1.A3. We find that a significant treatment effect in susceptible crimes is only evident for the 2004-2005 time period. This gives us further confidence that our estimate for this year is a unique event that cannot be likened to arbitrary fluctuations of previous years.

using data on tube journeys obtained from Transport for London. This records journey patterns for the main method of public transport around London and therefore provides a good proxy for shifts in the volume of activity around the city. We aggregate the journeys information to borough level and normalise it with respect to the number of tube stations in the borough.

Figure 1.5 shows how journeys changed year-on-year across the treatment and comparison groups. There is no evidence of a discontinuity in travel patterns corresponding exactly to the timing of the six week period of increased police presence. In fact the Figure shows a smoother change in tube usage, with the number of journeys trending back up and returning only gradually to pre-attack levels by the end of the year, but with no sharp discontinuity like the police and crime series.

Table 1.5 formally tests for a difference in journeys across treatment and comparison groups. It shows reduced form estimates using tube journeys as the dependent variable to test to what extent the fall in tube journeys after the attacks followed the pattern of the police deployment. The estimates indicate that total journeys fell by 22 percent (column 2, controls) over the period of Operation Theseus. However, some of this fall may have been due to a diversion of commuters onto other modes of public transport. This is particularly plausible given that two tube lines running through the treatment group were effectively closed down for approximately four weeks after 7th July. To examine this we instead normalize journeys by the number of *open* tube stations with the results reported in panel B of the Table. The effect is now smaller at 13 percent. Importantly, on timing, notice that reduced use of the tube persisted and carried on well after the police numbers had gone back to their original levels.

This final point about the *persistent* effect of the terror attacks on tube-related travel decisions is useful for illustrating the correlated shocks issue. As Table 1.5 shows, tube travel continued to be significantly lower in the treatment group for the whole period until the end of 2005. For example, columns (2) and (4) show that a persistent 10.3 percent fall in tube travel after the police deployment was completed, which is approximately half of the 21 percent effect seen in the Operation Theseus period. If the change in travel patterns induced by the terrorist attacks was responsible for reducing crime then we would expect some part of this effect to continue after the deployment.

At this point it is worth re-considering the week-by-week evidence presented in Figures 1.3a and 1.3b. A unique feature of the Operation Theseus deployment is that it provides us with two discontinuities in police presence, namely the way that the deployment was discretely switched on and off. The first discontinuity is of course related to the initial attack on July 7^{th}. Notably, along with an increased police deployment this first discontinuity is associated with a similarly timed shift in observable and unobservable factors. In particular, this first discontinuity in police deployment was also accompanied by a similarly acute shift in unobservable factors (that is, widespread changes in behaviors and attitudes towards public security risks – "panic" for shorthand). Because these two effects coincide exactly it is legitimate to raise the argument that the reduction in crime could have been partly driven by the shift in correlated unobservables.

However, the second discontinuity again provides a useful counterfactual. In this case the police deployment was "switched off" in an environment where unobservable factors were still in effect. Importantly, the Metropolitan Police never made an official public announcement that the police deployment was being significantly reduced. This decision therefore limits the scope for unobservable factors to explicitly follow or respond to the police deployment. It is therefore interesting to compare the treatment effect estimates immediately before and after the deployment was switched off. The estimated treatment interaction in week 85 (the last week of the police deployment) was -0.107 (0.043) while the same interaction in the two following weeks are estimated as being -0.040 (0.061) and -0.041 (0.045). This shows that crime in the treatment group increased again at the exact point that the police deployment was withdrawn. Furthermore, this discrete shift in deployment occurred as observable and unobservable factors that could have affected crime still strongly persisted (for example, recall the -10.3 percent gap in tube travel in Table 1.5 after the deployment was withdrawn).

More generally, this second discontinuity illustrates the point that any correlated, unobservable shocks affecting crime would need to be exactly and exquisitely timed to account for the drop in crime that occurred during Operation Theseus. Our argument then is that such timing is implausible given the decentralized nature of the decisions driving changes in unobservables. That is, the unobservable shocks are the result of individual decisions by millions of commuters and members of the public while Operation Theseus was a centrally determined policy with a clear “on” and “off” date. Indeed, the evidence on the police deployment that we show in this paper indicates that the Metropolitan Police’s response was quite deterministic. That is, deployment levels were raised in the treatment group while carefully keeping levels constant in the comparison group. Furthermore, police deployment levels were effectively restored to their pre-attack levels after Operation Theseus.[13] In contrast, shifts in travel patterns by inbound commuters did not match the timing and location of the police response.[14]

The issue of work travel decisions also uncovers a source of variation that we are able to exploit for evaluating the possible effect of observable, activity-related shocks. Specifically, any basic model of work and non-work travel decisions predicts interesting variations in terms of timing. For example, we would expect that faced with the terrorist risks associated with travel on public transport people would adjust their behavior differently for non-work travel. That is, the travel decision is less elastic for the travel to work decision compared to that for non-work travel. We would therefore expect that tube journeys would fall by proportionately more on weekends

[13] Our discussions with MPS policy officers indicate that big changes in the relative levels of ongoing police deployment in different boroughs occur only rarely. Relative levels of police deployment are determined mainly by centralised formulas (where the main criteria are borough characteristics) with changes determined by a centralised committee.

[14] More support for the hypothesis that changing travel patterns did not match the timing of changes in police presence follows from an analysis of Labour Force Survey (LFS) data. These data gives information on where people live and work enabling us to look at whether the number of inbound commuters to inner London changed. There is no evidence that the work travel decisions of people commuting in from outer London and the South East were affected by the attacks in that changes in the proportion of inbound commuters before and after the attacks are statistically insignificant, supporting the idea that modes of transport activity were affected more than travel volumes (see chapter Appendix Table 1.A4).

(when most non-work travel takes place) than on weekdays. This does seem to have been the case with tube journeys falling by 28 percent on weekends as compared to 20 percent on weekdays (see the lower panel of Table 1.5).

Thus there is an important source of intra-week variation in the shock to observables. If the shock to observables drives the fall in crime, then we would expect this to reflect a more pronounced effect of police on crime on weekends. Following this, we have re-estimated the baseline models excluding all observations relating to weekends.[15] This results in very similar coefficient estimates and only slightly larger standard errors, as shown in Table 1.6. Importantly, this means that our estimates are unaffected even when we drop the section of our crime data that is most vulnerable to the problem of correlated observable shocks.

A similar argument can be made in terms of correlated unobservable shocks. As already seen, there is a distinctive pattern to the timing of the fall in crime and its subsequent bounce back. For unobservable shocks to be driving our results their effect would have to be large and exquisitely timed to perfectly match the police and crime changes. However, basic survey evidence on risk attitudes amongst inner and outer London residents also suggests no significant difference in the types of attitudes that would drive a set of significant, differential unobservable shocks across treatment and comparison groups. Indeed, responses on attitudes to the terror attacks given by inner and outer London residents are closely comparable.[16] The attacks almost certainly had an impact on risk attitudes but they seem to be very similar in the treatment and comparison areas. From this we conclude that the effect of unobservables is likely to be minimal.

---

[15] Recall that our crime, police and tube journeys data are available at daily level for the years 2004-2005. This gives us the flexibility to drop Saturday and Sunday prior to aggregating to a weekly frequency.

[16] See chapter Appendix Table 1.A5.

1.46 Possible Crime Displacement

The final empirical issue we consider is that of crime displacement, both spatial and temporal. These two displacement effects have opposing effects on the crime-police relationship we have estimated. Firstly, spatial displacement into comparison areas is likely to impart a downward bias on our estimate as it will move criminal activity into the non-treated boroughs, increasing crime there and lowering the difference-in-difference estimate. Secondly, temporal displacement could impart an upward bias on our estimate. Criminals operating in the treatment group could delay their actions, thus contributing to a larger fall in crime during the policy-on period, but subsequently there will be a compensating increase in crime in the wake of the policy.

It turns out that we are not able to marshall evidence for either of these.[17] The paper by Draca, Machin and Witt (2011) looks at spatial displacement due to Operation Theseus in more detail. Here, we note the results of a robustness check where the comparison group is restricted to a set of adjacent and/or central London boroughs. If crime were displaced to these geographically closer boroughs, then we would see different estimates from the baseline estimates considered earlier. In particular, if crime rose in these nearby boroughs as a result of displacement then we would expect a smaller differences-in-differences estimate. Using these more matched boroughs (Adjacent and Central Ten[18]) produces very similar results to the earlier estimates. Compared to the baseline estimates discussed earlier (-0.132, with associated standard error 0.031), for susceptible crimes the estimated effects (standard errors) were -0.129 (0.040) for Adjacent and -0.108 (0.051) for Central Ten. Thus the estimates are similar, identifying a crime

[17] On the face of it, the lack of displacement for the kind of susceptible street-level crimes where we find effects might seem surprising. Unfortunately, we do not have data that would enable a stronger test to be undertaken (e.g. on tourist traffic in treatment and control boroughs, or on the attitudes of criminals).

[18] Adjacent boroughs were: Brent, Hackney, Hammersmith and Fulham, Lambeth, Newham, Southwark and Wandsworth. Central Ten boroughs were: Westminster, Camden, Islington, Kensington and Chelsea, Tower Hamlets (Treatment Group) and Brent, Hackney, Hammersmith and Fulham, Lambeth and Southwark.

fall of around 11-13 percent for susceptible crimes in central London relative to the (respective) comparison boroughs. In line with the earlier baseline results there was no impact on non-susceptible crimes. As such, it does not seem that contemporaneous spatial displacement occurred during Operation Theseus.[19]

Finally, the issue of temporal displacement in the treatment group can be looked at by referring back to the week-by-week estimates of treatment effects in Figure 1.3. As we have already noted, after the bounce-back in treatment borough crime that quickly occurred at the end of Operation Theseus, there is no evidence in the differenced models of any subsequent increase in treatment area crime. This seems to run against the hypothesis of inter-temporal substitution where criminal activity rebounds after the police deployment was withdrawn from the treatment boroughs.[20]

1.5 Conclusions

In this paper we provide new, highly robust evidence on the causal impact of police on crime. We find strong evidence that more police lead to reductions in what we refer to as susceptible crimes (i.e. those that are more likely to be prevented by police visibility including street crimes like robberies and thefts). Our starting point is the basic insight at the centre of Di Tella and Schargrodsky's (2004) paper, namely that terrorist attacks can induce exogenous variations in the allocation of police resources that can be used to estimate the causal impact of police on crime. Using the case of the July 2005 London terror attacks, our paper extends this strategy in two significant ways. First, the scale of the police deployment we consider is much

[19] As a further check for displacement effects, we also followed the approach of Grogger (2002) in contrasting crime trends between adjacent and non-adjacent comparison boroughs. However, again we could not uncover decisive evidence of between-borough displacement effects.

[20] Closer inspection of Figure 1.2(b) does show something of an upturn in crime in the comparison boroughs in the latter part of 2005. Indeed, it is possible that this could reflect a delayed spatial displacement effect. However, and counter to this, in the differenced statistical models treatment-control borough differences in the post-Operation Theseus period are not statistically significant (see the insignificant coefficients on T*Post-Attack 1.2 in Tables 1.2, 1.3, 1.4 and 1.6).

greater than the highly localized responses that have previously been studied. Together with the unique police hours data we use, this allows us to provide new, highly robust IV-based estimates of the crime-police elasticity. Furthermore, there is a novel *ceteris paribus* dimension to the London police deployment. By temporarily extending its resources (primarily through overtime) the police service was able to keep their force levels constant in the comparison group that we consider, while simultaneously increasing the police presence in the treatment group. This provides a clean setting to test the relationship between crime and police.

The research design delivers some striking results. There is clear evidence that the timing and location of falls in susceptible crimes closely coincide with the increase in police deployment. Crime rates quickly returned to pre-attack levels after the six week "policy-on" period. Shocks to observable activity (as measured by tube journey data) cannot account for the timing of the fall and it is hard to conceive of a pattern of unobservable shocks that could do so. However, as with other papers that adopt a 'quasi-experimental' approach, one might have concerns about the study's external validity. Using a very different approach from other papers looking at the causal impact of crime, our preferred IV causal estimate of the crime-police elasticity is approximately -0.38 (for susceptible crimes), which is strikingly similar to existing results in the literature (e.g. those of Levitt, 1997, Corman and Mocan, 2000, and Di Tella and Schargrodsky, 2004). Moreover, because of the scale of the deployment change and the very clear coincident timing in the crime fall, this elasticity is precisely estimated and supportive of the basic economic model of crime in which more police reduce criminal activity.

References

Becker, G and Rubinstein, Y (2004) "Fear and the Response to Terrorism: An Economic Analysis." University of Chicago mimeo.

Bertrand, M; Duflo, E and Mullainathan, S (2004) "How Much Should we Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119: 249-75.

Bloom, N (2009) "The Impact of Uncertainty Shocks." *Econometrica*, 77: 623-85.

Corman,H; and Mocan, N (2000). "A Time Series Analysis of Crime, Deterrence and Drug Abuse." *American Economic Review*, 87: 270-290.

Di Tella, R; and Schargrodsky, E (2004) "Do Police Reduce Crime? Estimate Using the Allocation of Police Forces After a Terrorist Attack." *American Economic Review*, 94: 115-133.

Donohue, J (1998) "Understanding the Time Path of Crime." *The Journal of Criminal Law and Criminology*, 88: 1423-1451.

Draca, M; Machin, S and Witt, R (2012) "Crime Displacement and Police Interventions: Evidence from London's "Operation Theseus." In *Crime, Institutions and Policies*, ed. Rafael Di Tella and Ernesto Schargrodsky. National Bureau of Economic Research Conference Volume.

Evans, W; and Owens, E (2007) "COPS and Crime." *Journal of Public Economics*, 91: 181-201.

Freeman, R (1999) "The Economics of Crime." In *Handbook of Labor Economics*, ed. Orley Ashenfelter and David Card, 3529-3572. Amsterdam: North Holland.

Gottfredson, M; and Hirschi, T (1990) *A General Theory of Crime*. Stanford: Stanford University Press.

Greater London Authority (GLA) Economics. 2005. *London's Economic Outlook*, 1-66.

Grogger, Jeffrey (2002) "The Effects of Civil Gang Injunctions on Reported Violent Crime: Evidence From Los Angeles County." *Journal of Law and Economics*, 45: 69-90.

Jacob, B; Lofgren, L; and Moretti, E (2007) "The Dynamics of Criminal Behavior: Evidence from Weather Shocks." *Journal of Human Resources*, 42: 489-527.

Klick, J; and Tabarrok, A (2005) "Using Terror Alert Levels to Estimate the Effect of Police on Crime." *The Journal of Law and Economics*, 48: 267-279.

Klockars, C (1983) *Thinking About Police*. New York: McGraw Hill.

Krueger, A (2006) *What Makes a Terrorist: Economics and the Roots of Terrorism (Lionel Robbins Lectures)*. Princeton: Princeton University Press.

Krueger, A; and Maleckova, J (2003) "Education, Poverty, Political Violence and Terrorism: Is there a Causal Connection?" *Journal of Economic Perspectives*, 17: 119-44.

Levitt, S (1997) "Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime." *American Economic Review*, 87: 270-290.

Levitt, S (2002) "Using Electoral Cycles in Police Hiring to Estimate the Effects of Police on Crime: Reply." *American Economic Review*, 92: 1244-50.

Machin, S; and Marie, O (2012) "Crime and Police Resources: The Street Crime Initiative." *Journal of the European Economic Association*.

McCrary, J (2002) "Using Electoral Cycles in Police Hiring to Estimate the Effects of Police on Crime: Comment." *American Economic Review*, 92: 1236-43.

OECD Economic Outlook (2002) *Economic Consequences of Terrorism*. Paris: OECD.

Sherman, L and Weisburd, D (1995) "General Deterrent Effects of Police Patrols in Crime “Hot Spots”: A Randomized Control Trial." *Justice Quarterly*, 12: 625-648.

1.6 Figures for Chapter 1.

FIGURE 1.1: MAP OF LONDON BOROUGHS

Notes: 32 boroughs of London. Treatment group for Operation Theseus police intervention includes: Camden, Kensington and Chelsea, Islington, Tower Hamlets and Westminster. See Table A1 of the Web Appendix for descriptive statistics on crime levels for the treatment and comparison groups.

FIGURE 1.2:

POLICE HOURS AND TOTAL CRIME (LEVELS) 2004-2005,

TREATMENT VERSUS COMPARISON GROUP

(a) Police Hours (per 1000 population)

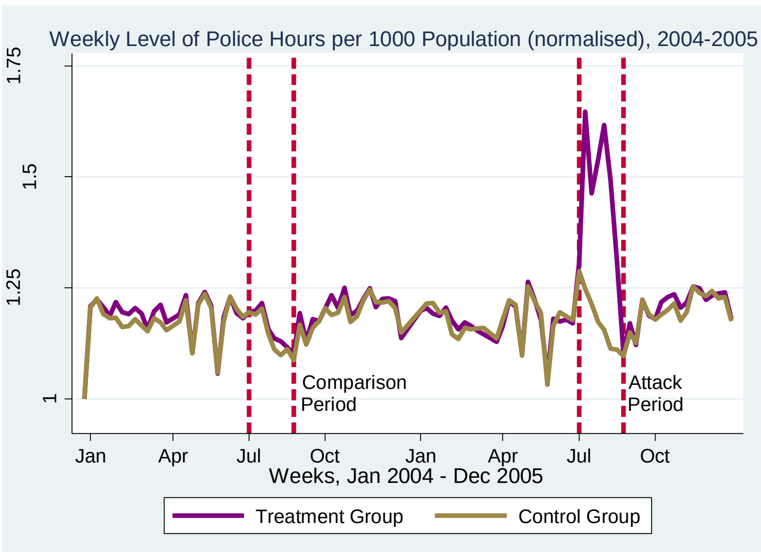

(b) Total Crimes (per 1000 population)

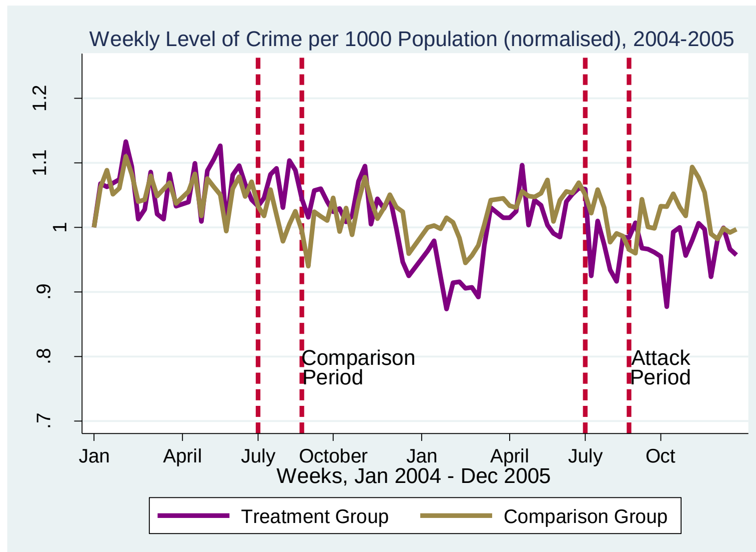

Notes: This Figure plots levels of police and crime for the treatment and comparison groups. Horizontal axis covers the period from January 2004 – January 2006. The values of police and crime have been normalised relative to the values in the first week of January 2004. Treatment and Comparison groups defined as per Figure 1.

FIGURE 1.3:

WEEK-BY-WEEK PLACEBO POLICY EFFECTS-POLICE HOURS AND SUSCEPTIBLE CRIMES

(a) Year-on-Year Change in Police Hours (per 1000 population) (b) Year-on-Year Change in Susceptible Crime Rate

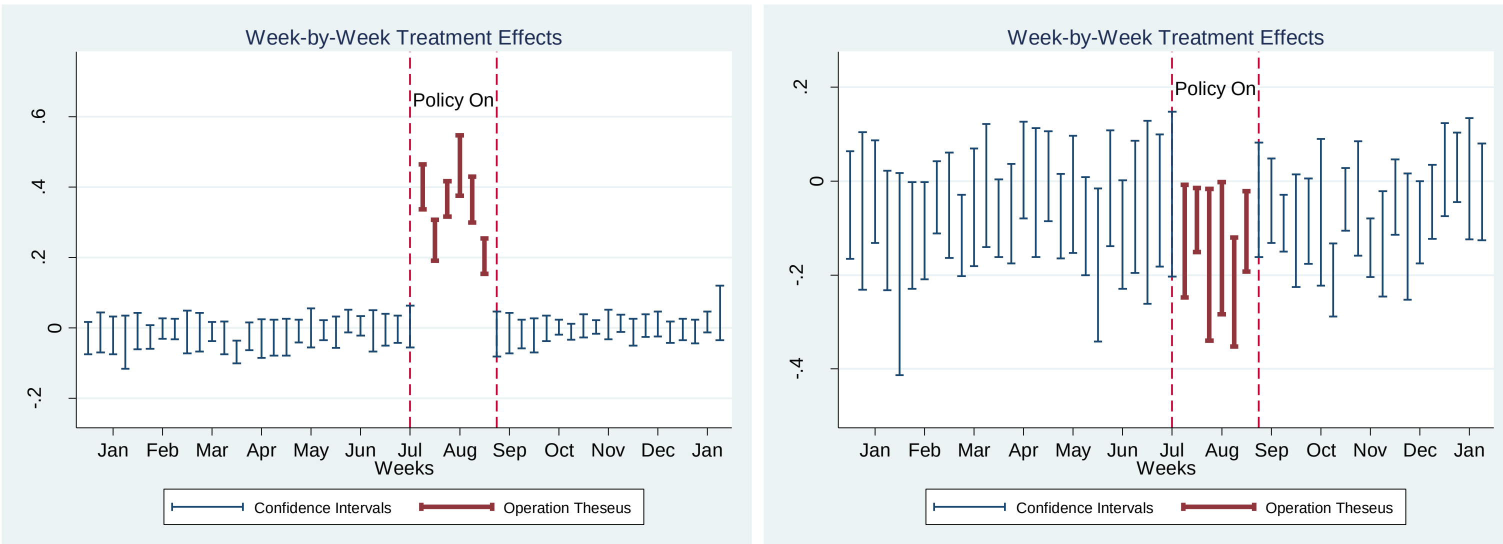

Notes: This Figure plots the coefficients and confidence intervals for week-by-week treatment*week interactions from January 2005-January 2006. These are estimated following the reduced form specifications in the main body of the paper. Standard errors clustered by borough. Note that since this is year-on-year, seasonally differenced data it reflects an underlying sample extending from January 2004-January 2006.

FIGURE 1.4:

YEAR-ON-YEAR CHANGES IN NUMBER OF TUBE JOURNEYS,

JANUARY 2004-JANUARY 2006.

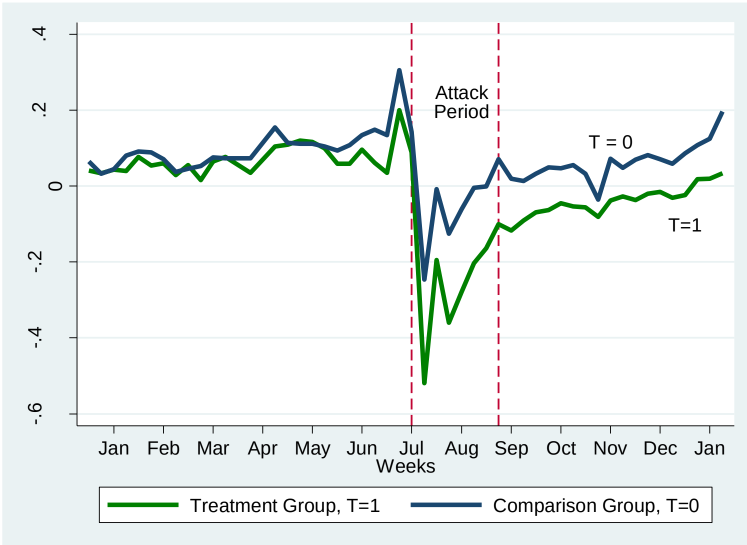

Notes: Horizontal axis covers the period from January 2005-January 2006. Note that since this is year-on-year, seasonally differenced data it reflects an underlying sample extending from January 2004-January 2006. The vertical axis measures the year-on-year log change in tube journeys. Tube journeys per station are measured as the sum of station entry and exit (i.e. inward and outward journeys) as recorded at station gates. Journeys per station are then aggregated to the borough and treatment/comparison group level for this graph. Data provided by Transport for London (TfL).

1.7 Tables for Chapter 1.

TABLE 1.1:

POLICE DEPLOYMENT AND MAJOR CRIMES, DIFFERENCES-IN-DIFFERENCES, 2004-2005

	(A) Police Deployment (Hours worked per 1000 Population)			*(B) Crime Rate (Crimes per 1000 Population)*		
	(1) Pre-Period	(2) Post-Period	(3) Difference (Post – Pre)	(4) Pre-Period	(5) Post-7/11	(6) Difference (Post – Pre)
T = 1	169.46	242.29	72.83	4.03	3.59	-0.44
T = 0	82.77	84.95	2.18	1.99	1.97	-0.02
Differences-in-Differences (Levels)			70.65*** (7.52)			-0.43** (0.16)
Differences-in-Differences (Logs)			0.35*** (0.04)			-0.11*** (0.04)

Notes: Post-period defined as the six weeks following 7/7/2005. Pre-period defined as the six weeks following 8/7/2004. Weeks defined in a Thursday-Wednesday interval throughout to ensure a clean pre and post split in the 2005 attack weeks. Treatment group (T = 1) defined as boroughs of Westminster, Camden, Islington, Tower Hamlets and Kensington-Chelsea. Comparison group (T = 0) defined as other boroughs of London. Police deployment defined as total weekly hours worked by police staff at borough-level. Standard errors are in parentheses.

TABLE 1.2:
DIFFERENCES-IN-DIFFERENCES REGRESSION ESTIMATES, POLICE DEPLOYMENT AND TOTAL CRIMES, 2004-2005.

	(A) Police Deployment (Hours Worked per 1000 Population)				(B) Total Crimes (Crimes per 1000 Population)				(C) OLS		(D) IV Estimates		
	Full	Split	+Controls	+Trends	Full	Split	+Controls	+Trends	Levels	Differences	Full	Split	+Trends
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
T*Post-Attack	0.081***				-0.052**								
	(0.010)				(0.021)								
T*Post-Attack1		0.341***	0.342***	0.356***		-0.111***	-0.109***	-0.056*					
		(0.028)	(0.029)	(0.027)		(0.027)	(0.027)	(0.030)					
T*Post-Attack2		-0.001	0.001	0.014		-0.033	-0.031	0.024					
		(0.011)	(0.010)	(0.016)		(0.027)	(0.028)	(0.054)					
ln(Police Hours)									0.785***				
									(0.053)				
Δln(Police Hours)										-0.031	-0.641**	-0.318***	-0.183***
										(0.051)	(0.301)	(0.093)	(0.066)
Controls	No	No	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	Yes	No	No	No	Yes	No	No	No	No	Yes
Number of Boroughs	32	32	32	32	32	32	32	32	32	32	32	32	
Number of Observations	1664	1664	1664	1664	1664	1664	1664	1664	3328	1664	1664	1664	1664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. Post-period for baseline models (1) and (5) defined as all weeks after 7/7/2005 until 31/12/2005 attack inclusive. Weeks defined in a Thursday-Wednesday interval throughout to ensure a clean pre and post split in the attack weeks. T*Post-Attack is then defined as interaction of treatment group with a dummy variable for the post-period. T*Post-Attack1 is defined as interaction of treatment group with a deployment "policy" dummy for weeks 1-6 following the July 7th 2005 attack. T*Post-Attack2 is defined as treatment group interaction for all weeks subsequent to the main Operation Theseus deployment. Treatment group defined as boroughs of Westminster, Camden, Islington, Tower Hamlets and Kensington-Chelsea. Police deployment defined as total weekly hours worked by all police staff at borough-level. Controls based on Quarterly Labour Force Survey (QLFS) data and include: borough unemployment rate, employment rate, males under 25 as proportion of population, and whites as proportion of population (following QLFS ethnic definitions).

TABLE 1.3:
TREATMENT EFFECTS BY MAJOR CRIME CATEGORY

	THEFTS, VIOLENCE AND SEX CRIMES					
Crime Category	Thefts		Violence		Sex Crimes	
	(1)	(2)	(3)	(4)	(5)	(6)
T*Post-Attack1	-0.139***	-0.082*	-0.124***	-0.108***	-0.078	-0.102
	(0.044)	(0.045)	(0.043)	(0.034)	(0.123)	(0.138)
T*Post-Attack2	-0.017	0.044	-0.054	-0.038	-0.080	-0.094
	(0.039)	(0.085)	(0.032)	(0.056)	(0.082)	(0.084)
Trends	No	Yes	No	Yes	No	Yes
Number of Boroughs	32	32	32	32	32	32
Number of Observations	1664	1664	1664	1664	1664	1664
	ROBBERY, BURGLARY AND CRIMINAL DAMAGE					
Crime Category	Robbery		Burglary		Criminal Damage	
	(1)	(2)	(3)	(4)	(5)	(6)
T*Post-Attack1	-0.131	-0.012	-0.035	-0.029	-0.047	-0.005
	(0.119)	(0.129)	(0.057)	(0.067)	(0.052)	(0.041)
T*Post-Attack2	-0.089	0.024	-0.093	-0.078	-0.018	0.020
	(0.098)	(0.149)	(0.059)	(0.075)	(0.043)	(0.057)
Trends	No	Yes	No	Yes	No	Yes
Number of Boroughs	32	32	32	32	32	32
Number of Observations	1664	1664	1664	1664	1664	1664

Notes: All specifications include week fixed effects. Standard clustered by borough in parentheses. Weighted by borough population. T*Post-Attack1 and T*Attack2 defined as per Table 2. Treatment group also defined as per Table 2. See Table A6 in the Web Appendix for definitions of the Major Crime categories in terms of the constituent Minor Crimes. Crime categories used follow the definitions provided by the Metropolitan Police Service (MPS).

TABLE 1.4:
SUSCEPTIBLE CRIMES VERSUS NON-SUSCEPTIBLE CRIMES, 2004-2005.

I. SUSCEPTIBLE CRIMES									
	(A) *Reduced Forms*				*(B)* *OLS*		*(C)* *IV Estimates*		
	Full	Split	+Controls	+Trends	Levels	Differences	Full	Split	+Trends
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
T*Post-Attack	-0.056**								
	(0.023)								
T*Post-Attack1		-0.131***	-0.132***	-0.067*					
		(0.031)	(0.031)	(0.035)					
T*Post-Attack2		-0.033	-0.033	0.033					
		(0.030)	(0.030)	(0.063)					
ln(Police Hours)					0.952***				
					(0.056)				
Δln(Police Hours)						-0.019	-0.694**	- 0.383***	-0.223***
						(0.063)	(0.336)	(0.105)	(0.074)
Controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	Yes	No	No	No	No	Yes
Number of Boroughs	32	32	32	32	32	32	32	32	32
Number of Observations	1664	1664	1664	1664	3328	1664	1664	1664	1664

Notes: All specifications include include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. Susceptible Crimes defined as: Violence Against the Person; Theft and Handling; Robbery. Non-Susceptible Crimes defined as: Burglary and Criminal Damage; Sexual Offences. Treatment group definitions and T*Post-Attack terms defined as per Table 2. Controls also defined as per Table 2.

TABLE 1.4:
SUSCEPTIBLE CRIMES VERSUS NON-SUSCEPTIBLE CRIMES, 2004-2005.
(continued)

(II) NON-SUSCEPTIBLE CRIMES

	(A) Reduced Forms				*(B) OLS*		*(C) IV Estimates*	
	Full	Split	+Controls	+Trends	Levels	Differences	Full	Split
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T*Post-Attack	-0.048*							
	(0.024)							
T*Post-Attack1		-0.033	-0.023	-0.015				
		(0.026)	(0.027)	(0.031)				
T*Post-Attack2		-0.053	-0.043	-0.033				
		(0.034)	(0.037)	(0.045)				
ln(Police Hours)					0.327***			
					(0.046)			
Δln(Police Hours)						-0.056	-0.597*	-0.065
						(0.094)	(0.337)	(0.078)
Controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	Yes	No	No	No	No
Number of Boroughs	32	32	32	32	32	32	32	32
Number of Observations	1664	1664	1664	1664	3328	1664	1664	1664

Notes: All specifications include include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. Susceptible Crimes defined as: Violence Against the Person; Theft and Handling; Robbery. Non-Susceptible Crimes defined as: Burglary and Criminal Damage; Sexual Offences. Treatment group definitions and T*Post-Attack terms defined as per Table 2. Controls also defined as per Table 2.

TABLE 1.5:
CHANGES IN TUBE JOURNEYS, BEFORE AND AFTER JULY 7TH 2005.

	log(Journeys/ Number of Stations)		*log(Journeys/ Number of Open Stations)*	
	(1)	(2)	(3)	(4)
	Baseline	Add Controls	Baseline	Add Controls
T*Post-Attack1	-0.212***	-0.215***	-0.133***	-0.137***
	(0.018)	(0.021)	(0.016)	(0.020)
T*Post-Attack2	-0.105***	-0.103***	-0.105***	-0.103***
	(0.008)	(0.008)	(0.008)	(0.008)
Controls	No	Yes	No	Yes
Observations	104	104	104	104
	log(Journeys/ Number of Stations)		*log(Journeys/ Number of Stations)*	
	(1)	(2)	(3)	(4)
	Weekdays	Weekends	Weekdays	Weekends
T*Post-Attack1	-0.196***	-0.281***	-0.197***	-0.294***
	(0.018)	(0.034)	(0.022)	(0.045)
T*Post-Attack2	-0.097***	-0.106***	-0.093***	-0.112***
	(0.010)	(0.032)	(0.010)	(0.030)
Controls	No	No	Yes	Yes
Observations	104	104	104	104

Notes: Borough level data collapsed by treatment and comparison group, 2 units over 52 weeks. All columns include week fixed effects. Standard errors clustered by treatment group unit in parentheses. All regressions weighted by treatment and comparison group populations. Results adjusted for closed stations (i.e. using the Number of Open Stations as denominator) do not count closed stations along the Piccadilly Line (Arnos Grove to Hyde Park Corner) and Hammersmith and City Line (closed from July 7th to August 2nd, 2005). Note that stations that intersect with other tube lines are not counted as part of this closure.

TABLE 1.6:
ESTIMATED CRIME TREATMENT EFFECTS WHEN EXCLUDING WEEKENDS

	(A) Susceptible Crimes		*(B) Non-Susceptible Crimes*	
	(1) Reduced Form	(2) IV	(3) Reduced Form	(4) IV
T*Post Attack1	-0.138***		0.005	
	(0.046)		(0.025)	
T*Post Attack2	-0.032		-0.037	
	(0.031)		(0.043)	
ln(Police Deployment)		-0.400***		0.017
		(0.150)		(0.072)
Controls	Yes	Yes	Yes	Yes
No of Boroughs	32	32	32	32
No of Observations	1664	1664	1664	1664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. These models estimate similar models to Table 4 but using a count of crimes per 1000 population that excludes all crimes occurring on weekends (i.e.: using only Monday-Friday). Treatment groups, T*Post-Attack terms and Crime Categories defined as in Table 4.

1.7 Appendix for Chapter 1

TABLE 1.A1: DISTRIBUTION OF CRIME IN LONDON BY MAJOR CATEGORY, 2004-2005

	(1) % of All Crimes	(2) Crime Rate (per 1000)	(3) % Occurring in Treatment Group	(4) Crime Rate in Treatment Group (per 1000)
(A) Susceptible Crimes				
Theft and Handling	44.0	106.1	28.0	234.0
Violence Against the Person	22.6	54.4	17.7	76.0
Robbery	4.6	11.1	15.5	13.5
(B) Non-Susceptible Crimes				
Burglary	12.3	29.6	17.4	40.5
Criminal Damage	15.5	37.4	13.6	40.0
Sexual Offences	1.1	2.7	21.8	4.6
Total	**100.0**	**241.1**	**21.3**	**204.2**

Notes: All major crimes occurring in the 32 boroughs of London between 1st January 2004 and 31st December 2005. Crime rate in column (2) calculated as number of crimes as per 1,000 members of population across all of London. Figures in columns (3) represent the proportion of crimes that take place in the treatment group by major category (for example, 28% of all Theft and Handling offences take place in the treatment group boroughs). Treatment group defined as boroughs of Westminster, Camden, Islington, Tower Hamlets and Kensington-Chelsea.

TABLE 1.A2: POLICE PATROLS AFTER JULY 7TH, 2005

Q: Have you seen more, less or about the same police patrols across London?	*Inner London*	*Outer London*
More (%)	70	62
About the Same (%)	20	27
Less (%)	5	3
Don't Know (%)	5	8
Total Respondents (Number)	248	361

Notes: Source is IPSOS MORI Survey. Exact wording of question: "Since the attacks in July, would you say you have seen more, less or about the same amount of police patrols across London?" Interviews conducted on 22-26 September 2005.

TABLE 1.A3: EXTENDED TIME PERIOD ANALYSIS BASED ON MONTHLY DATA, (BOROUGH LEVEL MODELS, DIFFERENCED ACROSS YEARS, 2001-2005)

Change in log(Susceptible Crimes Per 1000 Population)	(1)	(2)	(3)	(4)
Year on Year Changes	July/August 2001 – July/August 2002	July/August 2002 – July/August 2003	July/August 2003 – July/August 2004	July/August 2004 – July/August 2005
Treatment boroughs (T)	0.029	-0.060	-0.057	-0.098
Control boroughs (C)	0.072	-0.021	-0.028	0.007
T – C Gap	-0.043	-0.039	-0.029	-0.105***
	(0.029)	(0.030)	(0.043)	(0.031)

Notes: All models estimated in terms of seasonal differences (i.e. differenced relative to the same month in the previous year). Clustered standard errors in parentheses. Boroughs weighted by population. Treatment group defined as boroughs of Westminster, Camden, Islington, Tower Hamlets and Kensington-Chelsea. "Policy-on" period defined as July-August. Crime defined according to Susceptible category given in Table 4 of the main paper.

TABLE 1.A4: WORK TRAVEL PATTERNS INTO CENTRAL LONDON, BEFORE AND AFTER JULY 7TH

	(1) *Outer London Resident*	*(2)* *Rest of South-East Resident*
(A) Short-Run		
6 Week Before	0.166	0.035
6 Weeks After	0.175	0.037
Difference	0.005	0.002
	(0.022)	(0.008)
(B) Medium-Run		
12 Weeks Before	0.145	0.038
12 Weeks After	0.157	0.031
Difference	0.012	-0.006
	(0.021)	(0.005)
(C) Long-Run		
6 Months Before	0.155	0.034
6 Months After	0.160	0.031
Difference	0.005	-0.003
	(0.015)	(0.004)

Notes: Source is UK Quarterly Labour Force Survey (QLFS), 2004-2005. Standard errors clustered by week. Defined for all employed person aged 18-65 working in Central or Inner London. Column 1 defines all those residing in Outer London and working in Central or Inner London. Column 2 defines all those residing in the South East of England region and working in Central or Inner London.

TABLE 1.A5: SURVEY EVIDENCE ON COMMUNITY ATTITUDES, INNER VERSUS OUTER LONDON

Question & Response	(1) Inner London (%)	(2) Outer London(%)
(1) As a result of the attacks have you considered moving to live outside London or not?		
Yes	11	11
No	89	89
(2) How likely do you think it is London will experience another attack in the near future?		
Very likely	36	48
Somewhat likely	43	37
Not very likely	11	8
Not at all likely	4	3
Don’t Know	6	4
(3) As a result of the attacks, have you spent more or less time in Central London?		
More time	2	2
Less time	19	21
Made No Difference	78	76
(4) Since the July attacks have you personally or friends and relatives experienced any hostility on the basis of race or religion?		
Yes: Verbal Abuse	6	6
Yes: Physical Abuse	2	1
Yes: Felt Under Suspicion or Stared At	2	2
Yes: Generally Felt Hostility	2	2

Source: IPSOS MORI Survey.

TABLE 1.A6: LIST OF MINOR CRIMES BY MAJOR CATEGORY, 2004-2005.

Major Category	*Minor Category*	*As Proportion of Major Category Crimes (%)*
Violence Against Persons	Common Assault	31.5
	Harassment	21.4
	Aggravated Bodily Harm (ABH)	34.5
	Grievous Bodily Harm (GBH)	2.7
	Murder	0.1
	Offensive Weapon	4.0
	Other Violence	5.8
Sexual Offences	Rape	77.5
	Other Sexual	22.5
Theft and Handling	Picking Pockets	5.2
	Snatches	3.9
	Theft from Shops	10.5
	Theft / Taking of Pedal Cycles	5.3
	Theft / Taking of Motor Vehicles	12.6
	Motor Vehicle Interference and Tampering	0.1
	Theft from Motor Vehicles	23.7
	Other Theft	37.6
	Handling Stolen Goods	0.1
Robbery	Business Property	6.4
	Personal Property	93.6
Burglary	Burglary in a Dwelling	62.9
	Burglary in Other Buildings	37.1
Criminal Damage	Criminal Damage to Motor Vehicles	
	Criminal Damage to a Dwelling	44.3
	Criminal Damage to Other Buildings	28.7
	Other Criminal Damage	
		14.0
		13.0

Notes: All crimes occurring in the 32 boroughs of London between 1st January 2004 and 31st December 2005. Proportions calculated as share of the total crimes for the relevant Major Crime category. Both the Minor and Major Crime categories reported are those defined and used by police.

CHAPTER 2: REVOLVING DOOR LOBBYISTS

Abstract

Washington's 'revolving door' -the movement from government service into the lobbying industry- is regarded as a major concern for policy-making. We study how ex-government staffers benefit from the personal connections acquired during their public service. Lobbyists with experience in the office of a US Senator suffer a 24% drop in generated revenue when that Senator leaves office. The effect is immediate, discontinuous around the exit period and long-lasting. Consistent with the notion that lobbyists sell access to powerful politicians, the drop in revenue is increasing in the seniority of and committee assignments power held by the exiting politician.

Keywords: Lobbying, Revolving Door, US Congress, Political Connections, Political Elites.

JEL Classification: H11, J24, J45.

2.1 Introduction.

Lobbying is widely regarded as an important component of the US political sys- tem and has received considerable attention among scholars of political institutions and policy outcomes (Grossman and Helpman 2001; Ansolabehere, de Figueiredo, and Snyder 2003; Baumgartner et al. 2009).

One important characteristic of the US lobbying industry is the extent to which it is dominated by the "revolving door" phenomenon—i.e., the movement of federal public employees into the lobbying industry. For example, 56 percent of the revenue generated by private lobbying firms between 1998 and 2008 can be attributed to individuals with some type of federal government experience (see Table 2.1). Reflecting this, a recent ranking of the 50 top Washington lobbyists identified 34 as having federal government experience (Eisler 2007). Discussions of the revolving door often feature prominently in journalistic and watchdog group accounts of the lobbying industry (Attkisson 2008; Overby 2011), and unsurprisingly, regulation of the lobbying indus- try often devotes special attention to the revolving door phenomenon (Maskell 2010).

There are two main views regarding the importance of former government employees in the lobbying industry. The first view contends that revolving door lobbyists are valuable because "Washington is all about connections." In this view, experience in government allows former officials to develop a network of friends and colleagues that they can later exploit on behalf of their clients (Revolving Door Working Group 2005; Zeleny 2006). To illustrate, a recent profile of a top Washington lobbyist states that:

> (*Nancy*) *Taylor is a onetime health-policy director on Senator Orrin Hatch's Labor and Human Resources Committee, which had jurisdiction over much drug-patent*

legislation and food-and-drug laws. . . . Colleagues say as long as Hatch is in the Senate, Taylor will continue to bring in business (Eisler 2007).

A second view, often put forward by lobbyists themselves, is that the importance of individuals with prior government experience is due to higher innate ability Burger 2006; see also Diermeier, Keane, and Merlo 2005; and Mattozzi and Merlo 2008) and/or human capital accumulation (Heinz et al. 1993; Esterling 2004). The higher expertise of revolving door individuals can refer to policy matters, the inner workings of the legislative process, or even the preferences of particular constituencies. For example, a staffer-turned-lobbyist interviewed by the *Washington Post* argues that "[t]he technical processes of the House and Senate are not intuitive or widely known. Like with any service, people who have experience are going to be valuable to people who don't" (Eggen and Kindy 2009).

Evaluating the relative importance of political connections is therefore critical for understanding the value that lobbyists provide for their clients and, more generally, to assess the role of intermediaries in the lobbying process. It can also both contribute to our understanding of the incentives and selection issues facing government officials and help guide attempts to regulate the revolving door phenomenon.

In this paper we evaluate the extent to which ex–government officials convert their political contacts into lobbying revenue. We do this by studying how the lobbying revenue of congressional staffers–turned-lobbyists depends on the power of the congressional politicians for whom they have worked in the past. Ex–congressional staffers represent the largest single group of revolving door lobbyists (Table 2.1) and have been the focus of much of the popular discussion regarding the revolving door.

Our main finding is that lobbyists connected to US senators suffer an average 24 percent drop in generated revenue when their previous employer leaves the Senate. The decrease in revenue is out of line with preexisting trends, it is discontinuous around the period in which the connected senator exits Congress, and it persists in the long term. Measured in terms of median revenue per staffer-turned-lobbyist, this estimate indicates that the exit of a Senator leads to approximately a $182,000 per year fall in revenues for each affiliated lobbyist. We also find evidence that ex-staffers are less likely to work in the lobbying industry after their connected senators exit Congress.

We regard the above findings as evidence that connections to powerful, serving politicians are key determinants of the revenue that lobbyists generate. Consistent with this interpretation, we also find that the political power of the exiting politician is a good predictor of the drop in revenue suffered by the connected lobbyist. Lobbyists connected to exiting Senators who served in the Finance and Appropriations Committees and to representatives who served in the Ways and Means Committee suffer a substantial drop in revenue when the connected politician leaves office. Lobbyists connected to congressmen in neither of these powerful committees are statistically unaffected by their exits.

We interpret the connections that we study as relational capital (Burt 1992; Kale, Singh, and Perlmutter 2000): links of friendship, mutual trust, or even politician-specific knowledge that allow certain lobbyists to be more effective when particular politicians hold power. Of course, our results do not imply that lobbyists' general human capital is an irrelevant input for the lobbying production process. In fact, the best way to interpret our results is as estimates of the marginal effect of connections in this industry, with the other factors of production, such as ability and expertise, held constant and at sample levels. Nevertheless, the large magnitude of our estimates does indicate that connections to people in power represent a critical asset for the actors who serve as intermediaries in the lobbying process.

Studies on the congressional revolving door and on the personal relationships between lobbyists and congressmen are scarce, a surprising fact given the popular interest in, and policy relevance of, this topic. Early research used surveys of lobbyists to argue that policy and process knowledge is more important than personal connections (Salisbury et al. 1989; Heinz et al. 1993; Esterling 2004). Very recent evidence, using data made available thanks to the Lobbying Disclosure Act, emphasizes instead the role of personal connections. Eggers (2010) shows that revolving door lobbyists benefit from additional business when their affiliated party has control of the House or the Senate. Bertrand, Bombardini, and Trebbi (2011) measure connections using the contributions that lobbyists make to congressional election campaigns. They first show that the committee assignments of the congressmen who lobbyists are connected to represent a good predictor of the issues that lobbyists work on. More importantly, they also find that lobbyists switch issues in a predictable way as their connected congressmen switch committee assignments. Their conclusion that lobbyists' connections to politicians determine strongly what they do is consistent with the findings of this paper.

More generally, our paper is related both to the vast literature on the impact of money on politics (Ansolabehere, de Figueiredo, and Snyder 2003; Stratmann 2005) and to relatively recent research arguing that political connections matter for firm value (Fisman 2001; Johnson and Mitton 2003; Khwaja and Mian 2005; Knight 2006; Faccio 2006; Ferguson and Voth 2008). A remaining issue in the latter body of work is whether such connections can be traded. In other words, if connections to serving politicians are valuable assets, is there a market for them? Our findings suggest that the relation between clients and connected lobbyists in the US federal lobbying industry can be regarded as a market for political connections (arguably the largest in the world) in which companies or interest groups can acquire indirect links to serving politicians by hiring their former employees. Interestingly, this market appears to react quite rapidly to changing circumstances. For instance, we find below that the lobbying revenue generated by ex-staffers drops by a very large percentage one single semester after their ex-employers have left Congress.

The remainder of the paper is structured as follows. In Section 2.2 we present our data, in Section 2.3 we discuss our empirical strategy, and in Section 2.4 we discuss our main results. In Section 2.5, we conclude.

2.2 Data

The dataset used for this study is a lobbyist-level panel constructed from two main parts: a database of lobbying reports released under the Lobbying Disclosure Act of 1995 (hereafter, LDA) and a database of political employment that we construct from two new sources.

2.21 Lobbying database

The LDA required organizations to register and report information on their lobby- ing activities to the Senate Office of Public Records (SOPR). According to the act, lobbying activity is defined as contacts with officials, including background work performed to support these contacts. Two types of registrants are obliged to report under the LDA: lobbying firms and "self-filing" organizations that conduct in-house lobbying activities. The lobbying firm sector is comprised of firms who take on work for a number of different corporate and noncorporate clients. Self-filing organizations include corporations as well as peak industry groups and nonprofit single-issue organizations. Both types of registrants are required to report good-faith estimates of lobbying expenditures (for self-filing organizations) or lobbying revenue (for lobbying firms) every six months.

In this paper we focus on lobbyists working at lobbying firms[21]. The LDA defines a person as a "lobbyist" if they spend 20 percent or more of their time engaged in lobbying activities. Under the LDA, lobbying firms are required to file a separate report for each of their clients. The report must specify the revenue generated from that client, the issues for which the firm was engaged in lobbying, the house(s) of Congress and federal agencies contacted, and the names of the individual lobbyists serving that particular client during that period.

We use the version of the data compiled by the Center for Responsive Politics (CRP), a Washington-based nonprofit organization for the promotion of political transparency. Further details on how CRP has processed and compiled the SOPR informations are displayed in the online Appendix.

2.22 Political Employment

Our study utilizes two databases on the political employment and career histories of lobbyists. The first database is Lobbyist.info, a professional directory of lobbyists published by Columbia Books. This is an extensive lobbyist directory that contains contact information as well as career histories, biographical information, educa- tional background, and areas of expertise. From this, we extract information on lob- byists who have had periods of political

[21] The main reason for omitting in-house lobbyists from our analysis is the absence of meaningful individual-level productivity or earnings measures for this group of lobbyists. Expenditures of self-filing organizations include employee compensation, office overheads, and payments to vendors (which include but are not necessarily restricted to lobbying firms). Clearly, expenditures by an organization do not indicate whether a particular lobbyist is effective and/or well compensated at his or her job. For instance, it may be that the lower effectiveness of a lobbyist losing a connection translates into both a decrease in compensation *and* higher expenditures in other areas (such as outside vendors) in order to counteract this loss in effectiveness.

federal employment. The second database that we use is the Congressional Staffer Salaries (CSS) database. The CSS database is obtained by LegiStorm (a political information company) from published reports by the Secretary of the Senate and the Clerk of the House of Representatives[22].

We match names from each of the political employment databases against the lobbying reports data using a string-based algorithm. Numerous checks are made to ensure the accuracy of the match, with details reported in the online Appendix.

2.23. Descriptive Statistics

Table 1.1 gives some descriptive statistics of our dataset. We find in panel A that the average private lobbying firm employs 2.8 lobbyists and generates close to $700,000 in revenue.

Panel B reports information on the prevalence of former political employees. They represent 41.6 percent of all lobbyist-year observations. Over half of the group of former political employees is made up of former congressional staffers (22.6 percent of the total sample) while the remainder is comprised of ex-congressmen and former employees of government agencies, executive bodies, and presidential administrations. The focus of our study is the subgroup of former congressional staffers of politicians who served at some point in the 1998–2008 period covered by the LDA data.

Panel C reports the average revenue per lobbyist/year, for different types of private sector lobbyists. We calculate revenue per lobbyist in two alternative ways by summing what we

[22] The main information provided is: staffer name; start and end dates for a given employment spell; office of employment within Congress; the job title or position; and the total salary amount for a given job spell. We extract this information for all staffers working in personal and committee offices since the beginning of the database in 2000.

call “unweighted” and “weighted” revenues across lobbying con- tracts. For example, consider a $40,000 contract that is serviced by four lobbyists. The unweighted measure we define allocates each lobbyist an equal $40,000 in revenues from this contract. The weighted measure allocates $10,000 to each lobbyist. These revenues are then added up across all the contracts a lobbyist works on in a given period.

The two measures of revenue capture complementary aspects of the individual lobbyist–generated revenue. The unweighted measure essentially captures the revenue value of the “practice” with which each lobbyist is associated, since it aggregates the value of all the contracts in which an individual lobbyist is involved. Note that the revenue of a practice will typically be a subset of a lobbying firm revenue, especially if the firm is large. The weighted measure divides the value of each contract by the number of workers in it. It therefore captures the *revenue per worker* of the practice associated with an individual lobbyist.

The average weighted revenue per lobbyist/year ranges around $349,000 for the subgroup of congressional staffers we consider. This figure is closely in line with the reported salaries of lobbyists in this group. For example, the *Washington Post* reported in 2005 that “[s]tarting salaries have risen to about $300,000 a year for the best-connected aides eager to ‘move downtown from Capitol Hill’”. Industry news reports such as Brush (2010) also regularly use average revenue figures as a credible proxy for salary trends among Washington lobbyists. The average annual unweighted revenue per lobbyist takes much higher values. This is unsurprising since the full dollar value of a contract is assigned to each of the lobbyists involved in it. Figure 1 displays the distribution of unweighted revenue for ex-staffers and other lobbyists. Panel C also reveals that revolving door lobbyists generate significantly more revenue than other lobbyists. Lastly, note in panel D that revolving door lobbyists generate around 56 percent of total industry revenue. Out of this, more than half is accounted for by ex–congressional staffers.

2.3 Empirical Strategy

Our objective is to relate a measure of period-by-period revenues associated with each lobbyist to the number of distinct, *currently serving* politicians that the lobbyist has worked for prior to his entry into the lobbying industry. Our empirical model is as follows:

$$R_{it} = \alpha_i + \beta P_{it} + X'_{it}\theta + \gamma_t^{pc} + \epsilon_{it} \qquad (2.1)$$

where R_{it} is the (log) dollar revenue per individual lobbyist i in time period t. The vector X'_{it} represents time-varying characteristics measured at the individual level, α_{it} is the lobbyist-specific fixed effect, and γ_t^{pc} is a set of distinct time period effects for each subgroup of lobbyists. The time periods used are the 6-month periods required for reporting under the LDA, giving us 22 periods from 1998–2008, inclusive.

The key variable of interest is P_{it}, the count of currently serving politicians the lobbyist is linked to through his previous employment experience (note, however, that few lobbyists are connected to more than one senator or representative). There are two points worth highlighting here. Firstly, P_{it} only measures links with former political employers. Since we ignore the wider set of connections acquired by ex-staffers, we are probably undercounting the total value of political connections. Secondly, P_{it} is time-varying, as it goes down in value when a connected politician leaves office. The underlying hypothesis here is that politicians in office are particularly relevant for contemporary legislative outcomes. Serving politicians are able to vote on and influence the development of current legislation and this will be of interest to potential lobbying clients. The access that a lobbyist has with respect to his connected politician is therefore made obsolete when that politician is no longer in office.

Clearly, individual ability and expertise can be correlated with government experience as well as being a predictor of generated revenue. The inclusion of the α_i fixed effects implies, however, that β is identified from the variation in P_{it} described above. In other words, we compare the revenues of lobbyists who lose a political connection to the revenue of lobbyists whose connections remain constant.

We further narrow the comparison group by including γ_t^{pc} which are separate time effects for lobbyists connected to politicians in different parties (Democrat versus Republican) and chambers (House versus Senate). The inclusion of these time effects accounts for the fact that congressmen exits are likely to be correlated with shifts in party influence that can independently affect the ability to generate revenue. After including separate time dummies, our identifying assumption is that the revenue of lobbyists suffering a loss of a connection would have evolved similarly to the revenue of lobbyists connected to non-exiting politicians in the same party and chamber combination.

One relevant variation of equation (2.1) relates to timing. To study how lobbyists' revenues evolve in the individual periods just before and after the change in P_{it}, we can estimate:

$$R_{it} = \alpha_i + \sum_{l=-L}^{L} \beta_l P_{i(t_0+l)} + X_{it}^{i} + \gamma_t^{pc} + \epsilon_{it} \qquad (2.2)$$

where t_0 represents the transition period (i.e., when a politician exited Congress) and l flags the periods either before or after this period. This provides a set of time effects leading up to and following the transition period. We can use these, for instance, to examine whether revenue was already falling even before a connected politician's exit[23]. We can also study whether revenue spikes up in anticipation of an exit. Finally, equation (2.2) also allows us to study whether revenue recovers in the short or medium terms following the politician's exit.

Clearly, our identification strategy depends on the nature of the variation in P_{it}. Figures 2.2A and 2.2B show the number of lobbyists in the sample affected by the exit of a connected politician. In total, there are 257 lobbyists affected by these exits (94 for Senate exits, 163 for House exits), representing 20.9 percent of all ex-staffer lobbyists. Approximately half of exits are due to voluntary retirement of politicians. The next largest group of exits occurs as a result of defeats at reelection. The remainder of the exits is made up variously of lobbyists affected by politicians who die, leave due to a scandal, or run for another office (either successfully or unsuccessfully).

Finally, it should be noted that measurement error has the potential to attenuate our estimates in two ways. Firstly, there is the potential measurement error related to the name matching of lobbyists between our political employment and lobbying reports databases[24]. Secondly, there is measurement error related to R_{it} , arising from the fact that the size of the team serving a client is potentially an endogenous variable. For example, in a single-person firm it is straightforward to attribute revenues from clients to an individual lobbyist. But this becomes more complicated as the size of a firm increases, since as this happens team size becomes a firm

[23] This could be due, for instance, to the presence of "shared trends" between politicians and lobbyists. If low-ability lobbyists sort toward employment with low-ability politicians facing electoral defeat (and ability affects trends as well as levels), then revenue could be trending downward before exit.

[24] That is, lobbyists may have been either falsely matched to a politician or not assigned a true connection. It can be shown that this type of binary measurement error imparts a downward bias to β (Aigner 1973; Khwaja and Mian 2005).

choice variable. To minimize this problem, our regressions below use the unweighted measure of lobbyist revenues where we count the full value of contracts where a lobbyist is named and do not divide by team size before summing across a lobbyist contracts[25].

2.4 Main Results

2.41 Average Effects of Revolving Door Connections

Table 2.2 displays the estimates of empirical model (1). In column 1 we control only for individual lobbyist dummies. We find that being connected to a serving Senator is associated with 23 percent higher revenue, whereas the point estimate for a connection to a serving representative is only 9 percent and not statistically different from 0. Note that the difference in the estimated effects across the two chambers is consistent with the notion that it is the political power of the connected serving politician that matters. Senators are typically more powerful than representatives. For example, there are four times fewer senators than representatives and senators are uniquely able to wield filibuster powers that can slow down or completely block legislation.

---

[25] A second, more subtle, reason to use the unweighted measure is that it allows us to provide a better approximation to the marginal value of a political connection. To see this, consider a revolving door lobbyist A working with another lobbyist B. Imagine that together they generate $40,000 before the loss of A's connection and $30,000 after (and, obviously, that the loss in connection is orthogonal to other events affecting both lobbyists). While we cannot measure each individual's overall contribution to the team, we can reasonably conclude that the marginal value of A's connection was $10,000. This is what we would predict using the unweighted measure of revenue, while using the weighted measure we would instead estimate the value of A's connection as $5,000. Note, lastly, that we have estimated the full range of models reported in Section 2.5 using the weighted measure and have found very similar results. We report our main results using the weighted measure in the chapter Appendix.

As discussed earlier, other ex-staffers may not represent a valid comparison group for a lobbyist connected to an exiting politician. In the next three columns we successively narrow the comparison groups and we also control for lobbyist experience effects. In column 2 we add a full set a party-time dummies (allowing demand shocks to differ across former Republican and Democratic ex-staffers) while column 3 splits this further into party-chamber effects. Column 4 adds controls for lobbyist experience and its square. In this last and most comprehensive regression the exit of a connected senator is associated with 24 percent lower revenue.

Remarkably, the inclusion of extra controls only translates into very minor shifts in the coefficients for the senators and representatives variables. This suggests that politician exits are in practice a source of variation that is separate from party and chamber-related revenue shifts[26].

In the chapter Appendix we examine the robustness of our baseline results to the estimation of more stringent specifications. In particular, (i) we interact individual lobbyist dummies with the identity of the party in power; (ii) we include individual-specific time trends; and (iii) we add lobbying firm fixed effects to the individual lobbyist fixed effects. The point estimates are very consistent across specifications and the effect of being connected to a senator is always statistically significant at conventional levels.

Our estimate for connections to serving senators is economically as well as statistically significant. Evaluated at the mean of the yearly revenue generated by an ex-staffer's practice in our sample ($1,551,600 from Table 2.1), our estimate suggests that an active Senate connection

[26] In the chapter Appendix we explore whether lobbyists connected to the Democratic party earn more revenue in periods in which the Democrats control Congress. We find major revenue effects of party control of approximately 18%; see also Eggers (2010) for an estimate of party effects.

translates into approximately \$372, 000 per year. We believe, however, that the median, rather than the mean, value of revenue represents a better measure for the typical ex-staffer. The reason is that, as Figure 2.1 shows, the distribution of unweighted revenue has a very long right tail, with the median value being \$760,000, around half of the mean value. Evaluated at the median, our estimate suggests that an active Senate connection translates into approximately \$182,000 per year higher revenue for the value of an ex-staffer practice[27].

2.42 Timing Effects

In Table 2.2 we have presented evidence on the effect of political transitions on lobbyist revenue *averaged over time*. That is, we were comparing lobbyists' revenues in the average period before, and average period after, a given change in P_{it}. In Figure 2.3 we plot the results of estimating equation (2) for connections to serving senators. We use a window of 12 time periods (i.e., 6 years) around the time at which a politician's transition takes place. We have normalized the baseline to be period t_0, the last period in which a senator was still serving in Congress. The estimates should therefore be interpreted as relative to period t_0.

Several conclusions emerge from Figure 2.3. First, there is no strong evidence of either an upward or a downward trend in the periods leading up to a connected politician's exit. We can therefore reasonably rule out that our estimated average effects are due to the presence of shared trends between the fortunes of lobbyists and the politicians that they are connected to.

[27] What share of these \$182,000 reverts in terms of salary to the ex-staffer holding the connection is of course difficult to establish. Under the assumption that each of the lobbyists included in a contract gets rewarded according to the value of the assets that he contributes to the team, there would be a proportional loss in earnings for the individual ex-staffer.

Second, Figure 3 also seems inconsistent with the notion that abnormally high revenues occur prior to a politician's exit. It appears therefore that anticipation effects do not seem important, either because most exits are unanticipated or because lobbyists are unable or unwilling to extract higher revenue while a connection is still valuable. Third, there is also no evidence of reverse causality from lobbying revenue into the connected politician's exit. Note that period t_0 captures the last semester in which a politician served in Congress. We find that lobbying revenue during that semester, which could have potentially affected the politician's reelection chances, is very similar to that of previous semesters. It is only in the following semester, once the senator has already left office, that the connected lobbyist's revenue collapses. Thus, the timing of the drop in revenues relative to the timing of the politician's exit does not hint to the existence of reverse causality[28].

Our last conclusion from Figure 2.3 is that the negative effect of a connected politician's exit is highly persistent. There is evidence of a large drop in the period immediately after a politician's exit followed by some reversion. Lobbyists, however, are still subject to a 20 percent drop in revenues even 6 semesters after a politician's exit. This suggests both that lobbyists' links to their former employers are a major component of their overall political network and that lobbyists are not able to compensate for the loss of such a valuable connection using unobserved margins of adjustment.

---

[28] For politicians leaving at the end of their term, perhaps due to a reelection defeat, this comprises the period between July and December, which includes the November election date.

2.43 Effects Disaggregated by Political Power

Our interpretation of the average effects in Table 2.2 is that being connected to an individual holding political power allows a lobbyist to generate higher revenue. If our interpretation is correct we should expect individuals connected to more powerful politicians to suffer a larger drop in lobbying revenue when those politicians leave Congress.

One way to examine the hypothesis that it is political power that matters is to split politicians by their committee responsibilities at the legislator's point of exit. Ideally, we would like to create a different variable for connections to politicians in each of the different committees in the House and Senate. Our sample sizes, however, do not allow for such a level of disaggregation. We therefore decided to concentrate on what are arguably the two most important committees in the House and Senate—the Finance and Appropriations Committees in the Senate and the Ways and Means and Appropriations Committees in the House (Groseclose and Stewart 1998; Stewart and Groseclose 1999). These are committees with budget responsibilities and there- fore are particularly prone to be lobbied. These are also large committees that offer bigger cell sizes for our testing. We split lobbyists according to the service of their connected politicians on these committees at the time of the politicians' exits from Congress.

Table 2.3 displays the results. We find that lobbyists connected to Senators serving in the Finance and Appropriations Committees suffer losses in revenue of 36 percent and 45 percent, respectively, when those senators leave office. Similarly, we find that lobbyists connected to representatives serving on the Ways and Means Committee suffer losses in revenue of 35 percent when those representatives leave office. On the other hand, politicians serving in neither of these committees do not affect their affiliated lobbyists' revenue when they leave Congress[29].

As an additional exercise, we also studied whether there is evidence of an increase in generated revenue when connected politicians remain in Congress and join important committees. We found that being connected to senators in the Finance Committee and Representatives in the Ways and Means Committee is indeed associated with significant revenue premiums. We regard this evidence as consistent with the main message of the paper, and we display it in the online Appendix.

2.44. Participation in the Lobbying Industry

The models and estimates presented above show a strong effect of changes in political connections on lobbyist revenues. As a result of this revenue effect, changes in political connections could also affect an ex-staffer's decision about whether to work in the lobbying industry at all. To study whether this is the case, we expand our dataset to include, for each individual, every period following the end of their employment as a staffer[30]. We then define the variable A_{it} = 1 if individual *i* served any client during period *t* and 0 otherwise. Our new dataset contains 16,882 observations and the mean of A_{it} is 0.62. We then estimate a variation of equation (2.1) using A_{it} as our new dependent variable.

---

[29] In the Senate, the "neither" group is statistically different from the "Finance" and "Appropriations" groups at the 1 percent and 10 percent levels, respectively. In the House, the "neither" group is statistically different from the "Ways and Means" group at the 1 percent level.

[30] For instance, if an individual left his or her job in Congress at the end of 2002 then our sample contains observations for this individual over the period 2002–2008, whether or not he or she was actually working as a lobbyist.

Table 2.4 displays the results of estimating our variation of equation (1) using linear probability models. We find that being connected to a currently serving senator is associated with 27 percent higher likelihood of working as a lobbyist. Our estimate for the representative effect is much smaller and statistically insignificant. Again, including or excluding lobbyist experience and separate subgroup time effects has little impact on our coefficients.

Our findings from Table 2.4 are robust to the use of nonlinear (logit) models. We also find very similar effects when we expand our dataset further to include, for every lobbyist, each period between 1998 and 2008. This evidence can be found in the online Appendix[31].

2.5 Concluding Remarks

In this paper we show that ex–government officials extract monetary rents in terms of generated lobbying revenue from their personal connections to elected representatives. Overall, our findings suggest that access to serving officials is a scarce asset that commands a premium in the market for lobbying services.

Before we discuss the contributions to existing research, it is worth pointing out some limitations of our study. While our focus on ex-staffers provides us with a strong identification strategy, the extent to which our findings apply to other lobbyists can be debated. Some revolving door lobbyists (e.g., ex-congressmen) may benefit even more from connections to

[31] The existence of a participation effect suggests that we may be underestimating the value of connections in our main regressions. This is because we would expect the hardest-hit lobbyists to be more likely to exit the industry. To study this, we assigned the exiting lobbyists revenue observations equal to their final observation before dropping out of the industry. We regard this final revenue figure as an upper bound for what the exiting lobbyist would have earned had he decided to remain in the industry. The effect of applying this bound is shown in the Chapter Appendix.

serving government officials, while others (e.g., ex–agency staffers) may rely less on connections and more on the expertise acquired while in government. Obviously, our estimates are not easily extrapolated to lobbyists with no government experience, although they help to explain the fact that these lobbyists generate substantially less revenue (Table 2.1) and are known to command lower salaries (Brush 2010). While acknowledging our exclusive focus on ex-staffers, we note that these represent a leading fraction of the lobbying industry, accounting for 34 percent of total revenue (Table 2.1).

A related limitation stems from our unique focus on the relationship between ex-staffers and their former employers. We would expect staffers to have developed a wide range of relationships with both elected representatives and other staffers. Since our study is based on one single relationship (albeit an important one), it probably undercounts the value of political connections.

Lastly, we provide no direct evidence on the existence of a "payback" for lobbying clients. Our contribution to the vast literature estimating the returns to lobbying and campaign contributions (de Figueiredo and Silverman 2006) is therefore only indirect. Namely, the fact that firms and interest groups are eager to hire the services of well-connected individuals suggests that they expect a return in terms of favorable legislative outcomes. Likewise, our findings cannot discriminate between alternative theories of what lobbying is. Our connected lobbyists could arguably serve as a conduit of both quid pro quo offers (Grossman and Helpman 1994) and information (Austen-Smith 1996). That said, existing theories do not account for the role of intermediaries in the lobbying process and this is an area where our findings suggest that future research could provide valuable insights.

With the caveats above, we believe that our findings have implications in terms of the career incentives of staffers (and probably other government officials). We have shown that staffers' political connections are a perishable asset; in other words, they last only as long as the connected politicians remain holding office. This implies that staffers may have relatively short careers. Once a connection to a powerful senator has been established, it may make sense to move into lobbying and cash in on this unique asset while it is still valuable. Of course, the existence of rents associated with post-government employment could widen the pool of applicants for staffer positions, and potentially allow congressmen to hire high-ability individuals at the lower salaries that the public sector typically offers (Caselli and Morelli 2004; Besley 2005).

Our paper also has the potential to inform policy. One common instrument to regulate the revolving door phenomenon is to impose "cooling off" periods to officials leaving public office (Ethics Reform Act of 1989; Honest Leadership and Open Government Act of 2007; for a review, see Maskell 2010). The perishable nature of ex-staffers' assets suggests that such restrictions could in fact be quite useful to a legislator interested in significantly decreasing the attractiveness of a lobbying career for ex–government officials.

Finally, this paper contributes to our more general understanding of what makes workers valuable in professional services industries. The empirical results show that professional experts working in the lobbying industry are valued not only for their technical knowledge but in large part also for their personal connections (Oyer and Schaefer 2010; Singh, Hansen, and Podolny 2010). Hence, one insight from our study is that a large proportion of the premia that experts command in professional industries is likely to be comprised of so-called relational capital (Burt 1992; Kale, Singh, and Perlmutter 2000). Furthermore, the relational capital that this paper has studied is clearly valuable only in a very specific geographic and sectoral setting. A second insight from our paper is therefore that relational capital is likely to represent a large part of what is usually classified as industry-specific human capital (Neal 1995).

REFERENCE LIST

Aigner, DJ (1973) “Regression With a Binary Independent Variable Subject to Errors of Observa- tion.” *Journal of Econometrics* 1 (1): 49–59.

Ansolabehere, S; de Figueiredo, JM and Snyder, Jr, JM (2003) “Why Is There So Lit- tle Money in US Politics?” *Journal of Economic Perspectives* 17 (1): 105–30.

Attkisson, S (2008) “The Revolving Door: Lawmaker to Lobbyist.” *CBS News*, May 9

Austen-Smith, D. 1996. “Interest Groups: Money, Information, and Influence.” In *Perspectives on Public Choice: A Handbook*, edited by Dennis C. Mueller, 296–321. New York: Cambridge Uni- versity Press.

Baumgartner, FR, Berry, JM; Hojnacki, M; Kimball, DC and Leech, BL (2009) *Lobbying and Policy Change*. Chicago: University of Chicago Press.

Bertrand, M; Bombardini, M; and Trebbi, F (2011) “Is It Whom You Know or What You Know? An Empirical Assessment of the Lobbying Process.” NBER Working Paper 16765.

Brush, S (2010) “Democrats’ Ambitious Legislative Agenda Pushes K Street Salaries Skyward.” *The Hill*, February 23.

Burger, Timothy J (2006) “The Lobbying Game: Why the Revolving Door Won’t Close.” *Time Magazine*, February 16.

Burt, Ronald S (1992) *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Caselli, F; and Morelli, M (2004) “Bad Politicians.” *Journal of Public Economics* 88(3–4): 759–82.

de Figueiredo, JM; and Silverman, BS (2006) “Academic Earmarks and the Returns to Lobbying.” *Journal of Law and Economics* 49 (2): 597–625.

Diermeier, D; Keane, M; and Merlo, A (2005) “A Political Economy Model of Congressional Careers.” *American Economic Review* 95 (1): 347–73.

Eggen, D; and Kindy, K (2009) “Familiar Players in Health Bill Lobbying.” *Washington Post*, July 6.

Eggers, AC (2010) “The Partisan Revolving Door.” Unpublished.

Eisler, K (2007) “Hired Guns: The City’s 50 Top Lobbyists.” *The Washingtonian Magazine*, June 1.

Esterling, KM (2004) *The Political Economy of Expertise: Information and Efficiency in American National Politics*. Ann Arbor: University of Michigan Press.

Faccio, M (2006) "Politically Connected Firms." *American Economic Review* 96 (1): 369–86.

Ferguson, T; and Voth, H-J (2008) "Betting on Hitler—The Value of Political Connections in Nazi Germany." *Quarterly Journal of Economics* 123 (1): 101–37.

Fisman, R (2001) "Estimating the Value of Political Connections." *American Economic Review* 91 (4): 1095–102.

Groseclose, Tim, and Charles Stewart III. 1998. "The Value of Committee Seats in the House, 1947–91." *American Journal of Political Science* 42 (2): 453–74.

Grossman, GM; and Helpman, E (1994) "Protection for Sale." *American Economic Review* 84 (4): 833–50.

Grossman, GM; and Helpman, E (2001) *Special Interest Politics*. Cambridge, MA: MIT Press.

Heinz, JP; Laumann, EO; Nelson, RL; and Salisbury, RH (1993) *The HollowCore: Private Interests in National Policy Making*. Cambridge, MA: Harvard University Press.

Johnson, S; and Mitton, T (2003) "Cronyism and Capital Controls: Evidence from Malaysia."*Journal of Financial Economics* 67 (2): 351–82.

Kale, P; Singh, H; and Perlmutter, H (2000) "Learning and Protection of Proprietary Assets in Strategic Alliances: Building Relational Capital." *Strategic Management Journal* 21(3): 217–37.

Khwaja, A; and Mian, A (2005) "Do Lenders Favor Politically Connected Firms? Rent Provision in an Emerging Financial Market." *Quarterly Journal of Economics* 120 (4): 1371–411.

Knight, B (2006) "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush Gore 2000 Presidential Election." *Journal of Public Economics* 91 (1–2): 389–409.

Maskell, J (2010) "Post-Employment, 'Revolving Door,' Laws for Federal Personnel." Research Service Report for Congress. Washington, DC: Congressional Research Service.

Mattozzi, A; and Merlo, A (2008) "Political Careers or Career Politicians?" *Journal of Public Economics* 92 (3–4): 597–608.

Neal, D (1995) "Industry-Specific Human Capital: Evidence from Displaced Workers." *Journal of Labor Economics* 13 (4): 653–77.

Overby, P (2011) "Outgoing FCC Commissioner to Lobby for Comcast." *National Public Radio*, May 12.

Oyer, P; and Schaefer, S (2010) "Firm/Employee Matching: An Industry Study of American Lawyers." Unpublished.

Revolving Door Working Group (2005) "A Matter of Trust: How the Revolving Door

Undermines Public Confidence in Government—And What To Do About It”. http://www.cleanupwashington.org/ documents/RevovDoor.pdf.

Salisbury, RH; Johnson, P; Heinz, JP; Laumann, EO and Nelson, RL (1989) “Who You Know versus What You Know: The Uses of Government Experience for Washington Lobbyists.” *American Journal of Political Science* 33 (1): 175–95.

Singh, J; Hansen, MT; and Podolny, JM (2010) “The World Is Not Small for Everyone: Inequity in Searching for Knowledge in Organizations.” *Management Science* 56 (9): 1415–38.

Stewart, C; and Groseclose, T (1999) “The Value of Committee Seats in the United State Senate, 1947–91.” *American Journal of Political Science* 43 (3): 963–73.

Stewart, C; and Woon, J (2009) “Congressional Committee Assignments, 103rd to110th Congresses, 1993–2008: Senate and House.” Unpublished.

Stratmann, T (2005) “Some Talk: Money in Politics: A (Partial) Review of the Literature.” *Public Choice* 124 (1–2): 135–56.

Zeleny, J (2006) “As Guard Changes in Congress, Lobbyists Scramble.” *New York Times*, November 15.

2.6 Tables for Chapter 2.

TABLE 2.1: DESCRIPTIVE STATISTICS: US FEDERAL LOBBYING FIRMS, 1998-2008.

Panel A – By Lobbying Firms	
Mean Number of Lobbyists	2.8
Total Revenue/Expenditures	687.8
Total Number of Firms/Organizations	3,960
Panel B - Types of Lobbyists	
Revolving Door Lobbyist	0.416
Ex-Congressman	0.029
Ex-staffer:	0.226
- of politician serving pre-1998	0.064
- of politician serving post-1998	0.134
- of a congressional committee	0.027
Outside Congress	0.162
Panel C - Mean Revenue or Expenditure	
Weighted	
Revolving Door Lobbyists	309.9
-Ex-Congressmen	339.6
-Ex-staffers:	349.7
-Outside Congress	253.6
Other lobbyists	170.0
Unweighted	
Revolving Door Lobbyist	1355.5
-Ex-Congressmen	1287.8
-Ex-staffers:	1551.6
-Outside Congress	1109.8
Other lobbyists	682.8
Panel D - Share of total industry	
Revenue, by type of lobbyist	
Revolving Door Lobbyist	0.559
-Ex-Congressmen	0.043
-Ex-staffers:	0.343
-Outside Congress	0.182
Number of Lobbyists (Total)	15,315
Number of Lobbyist-Period Observations	98,705

Note: Panel(A) based on 1998-2008 panel. Panels (B) and (C) based on 1998-2008 lobbyist-level panel. Panel (C) presents annualised measures of revenue or expenditure per lobbyist. Panel (D) aggregates the weighted revenues of lobbyists in order to calculate revenue shares by type. 'Ex-Congressman' denotes former members of the House or Senate who are lobbyists. 'Ex-staffer' represents lobbyists who have employment experience as Congressional staffers.

TABLE 2.2: AVERAGE EFFECTS OF REVOLVING DOOR CONNECTIONS ON LOBBYING REVENUE.

	Dependent Variable: (log) revenue per lobbyist			
		Plus Party	Plus Chamber	Plus experience
	(1)	(2)	(3)	(4)
Number of Senators	0.23***	0.23***	0.21***	0.24***
	(0.07)	(0.07)	(0.07)	(0.07)
Number of Representatives	0.09*	0.07	0.08	0.10*
	(0.05)	(0.05)	(0.05)	(0.05)
Individual dummies	Yes	Yes	Yes	Yes
Time	Yes	No	No	No
Time*Party	No	Yes	No	No
Time*Party*Chamber	No	No	No	Yes
Lobbyist Experience	No	No	No	Yes
Individuals	1,113	1,113	1,113	1,113
Observations	10,418	10,418	10,418	10,418

Notes: This table presents the average effects of political connections on ex-staffers lobbying revenue. The dependent variable is the log of the revenue generated from all the clients that an individual lobbyist serves in a time (semester) period. The two main independent variables are the number of senators and representatives that an individual lobbyist worked for previous to entering the lobbying industry *and are serving in Congress in that time period*. All regressions use a sample containing ex-staffers-turned-lobbyists and include both individual lobbyist dummies and time effects (i.e., semester dummies). Column 2 allows for different time effects for lobbyists connected to politicians in different parties (i.e., Democrats versus Republicans). Columns 3 and 4 allow for different time effects for lobbyists connected to politicians in different party/chamber combinations (i.e., Democrats in the Senate, etc.). Column 4 includes lobbyist experience (i.e., number of periods that a lobbyist appears in the sample) in quadratic form. Standard errors are clustered by lobbyist. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

TABLE 2.3: EFFECTS DISAGGREGATED BY POLITICIAN COMMITTEE ASSIGNMENTS.

	Dependent Variable: (log) revenue per lobbyist	
		Plus party-chamber and experience
	(1)	(2)
Number of Senators		
in Finance	0.36***	0.36***
	(0.09)	(0.09)
in Appropriations	0.40***	0.45***
	(0.16)	(0.15)
in Neither	-0.12	-0.11
	(0.12)	(0.12)
Number of Representatives		
in Ways and Means	0.37***	0.35***
	(0.10)	(0.10)
in Appropriations	0.07	0.06
	(0.11)	(0.11)
in neither	-0.01	0.03
	(0.06)	(0.06)
Individual dummies	Yes	Yes
Time	Yes	No
Time*Party*Chamber	No	Yes
Lobbyist Experience	No	Yes
Individuals	1,113	1,113
Observations	10,418	10,418

Notes: This table presents the effects of Table 2 separately for different levels of politician committee assignments. The dependent variable is as in Table 2. The main independent variables are as in Table 2, with the exception that connections to senators and representatives are disaggregated by the politician committee assignments *at the time of leaving Congress*. All regressions use a sample containing ex-staffers-turned-lobbyists and include individual lobbyists dummies and time effects (i.e., semester dummies). Column 2 allows for different time effects for lobbyists connected to politicians in different party-chamber combina- tions (i.e., Democrats in the Senate, etc.) and also includes lobbyist experience (i.e., number of periods that a lobbyist appears in the sample) in quadratic form. Standard errors are clustered by lobbyist.

***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

TABLE 2.4 – PARTICIPATION IN THE LOBBYING INDUSTRY.

	Dependent variable: A_{it} = 1 if individual *i* generated positive revenue in period t	
		Plus party-chamber and experience
	(1)	(2)
Number of Senators	0.19***	0.27***
	(0.05)	(0.05)
Number of Representatives	0.02	0.06
	(0.04)	(0.04)
Individual dummies	Yes	Yes
Time	Yes	No
Time*Party*Chamber	No	Yes
Lobbyist Experience	No	Yes
Individuals	1,113	1,113
Observations	16,882	16,882

Notes: This table presents the effects of political connections on ex-staffers' participation in the lobbying industry. The dataset contains, for each individual, every period following the end of their employment as a staffer. The dependent variable takes value 1 when an individual generates positive lobbying revenue in a time (semester) period and 0 otherwise. For instance, if an individual left his or her job in Congress at the end of 2002, then our sample contains observations for this individual over the period 2002–2008, whether or not he or she was actu- ally working as a lobbyist. The main independent variables are as in Table 2. All regressions include individual lobbyists dummies and time effects (i.e., semester dummies). Column 2 allows for different time effects for lobbyists connected to politicians in different party/chamber combinations (i.e., Democrats in the Senate, etc.) and also includes lobbyist experience (i.e., number of periods that a lobbyist appears in the sample) in quadratic form. Standard errors are clustered by lobbyist.

***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

2.7 Figures for Chapter 2

FIGURE 2.1

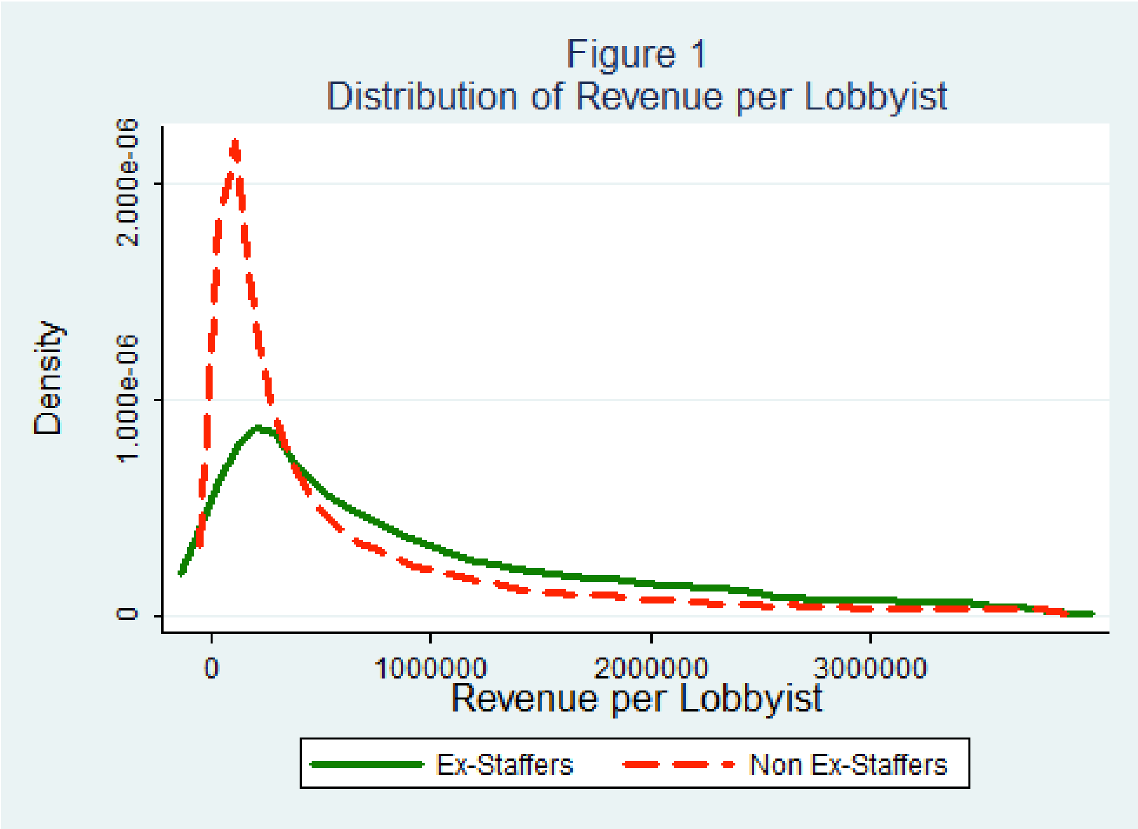

<u>FIGURE 2.2: REASONS FOR LOSS OF SENATE AND HOUSE CONNECTIONS, BY NUMBERS OF LOBBYISTS.</u>

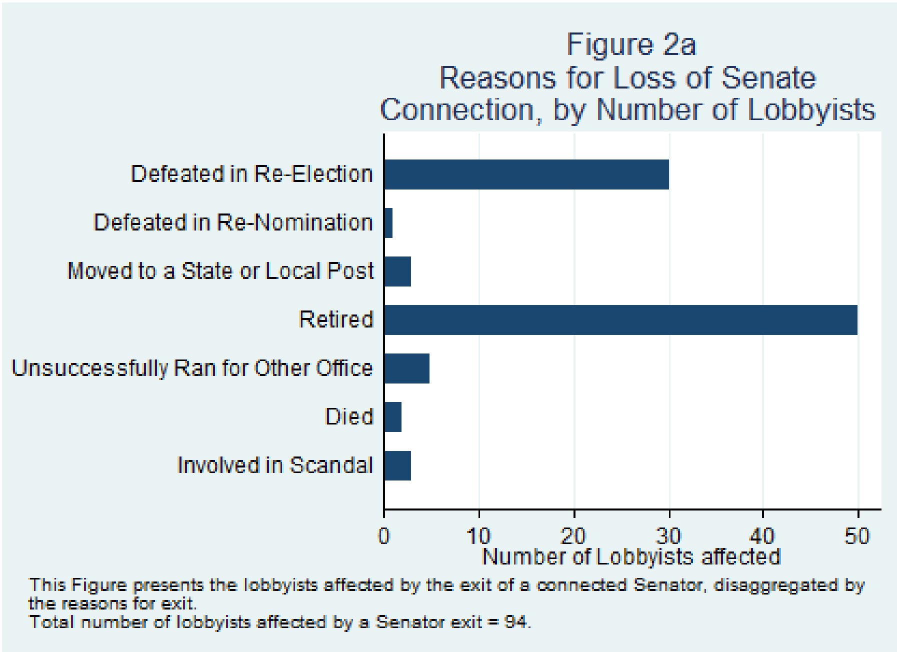

Figure 2b
Reasons for Loss of House Connection, by Number of Lobbyists

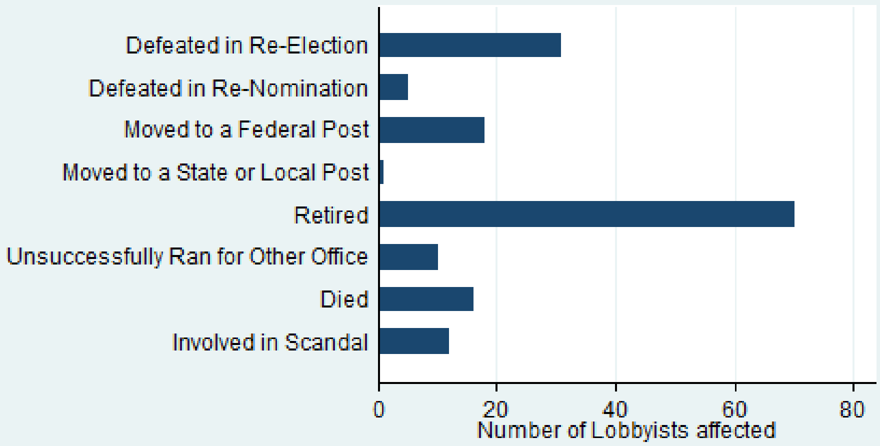

This Figure presents the lobbyists affected by the exit of a connected Representative, disaggregated the reasons for exit.
Total number of lobbyists affected by a Representative exit = 163.

FIGURE 3: TIMING EFFECTS AND THE LOSS OF CONNECTIONS.

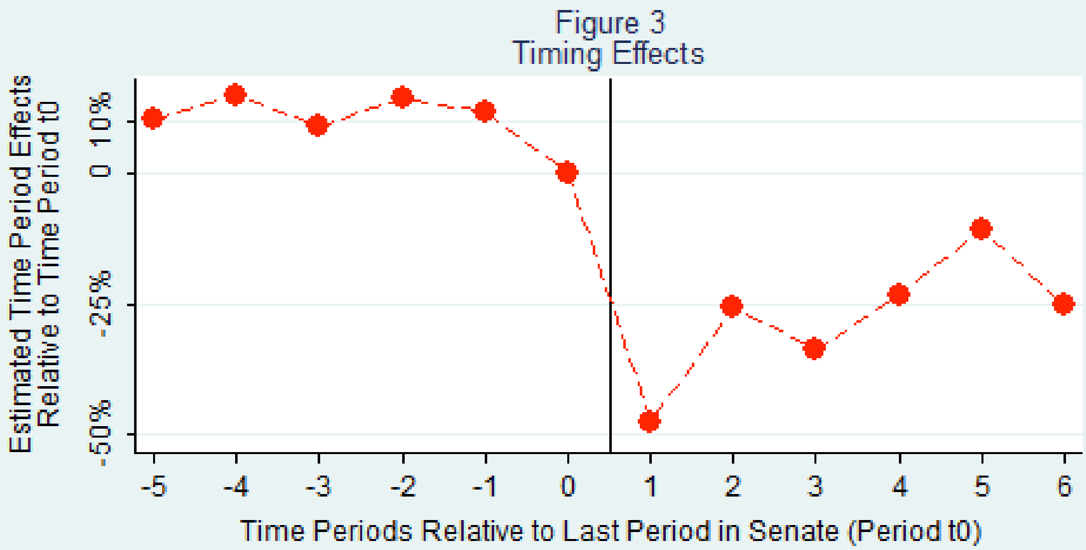

Note: Figure 3 displays the estimated time period effects leading up to and following the transition (exit from Senate) period from equation (2) in Section 3. Each period comprises six months. The left hand side is lobbyist (log) revenue as in Table 2. The sample includes all lobbyists with at least 4 observed periods before a transition and at least 2 periods after. Right-hand side variables include individual lobbyist dummies and party-chamber time dummies, as well as estimated period effects. Every estimated effect is relative to period t0 - the last period in which a Senator is still serving in Congress. Since period t0 is the omitted group, there is no estimated effect for this group. We however display an 'effect' of 0 at t0 to aid visual analysis.

2.8 Appendix for Chapter 2

In this appendix we provide additional information on various segments of the data used to construct the lobbyist-level panel used in the paper. We also provide more details on the name matching procedure used to link the databases on lobbying reports and political employment.

Lobbying Reports

We use the LDA data sourced from the Senate Office of Public Records (SOPR) and compiled by the Centre for Responsive Politics (CRP) as part of their 'Open Secrets' database. The CRP is a non-profit and non-partisan organization with a stated mission of collating information on all types of politically related expenditures (i.e. campaign contributions, lobbying expenditures, member personal finances, etc.) and facilitating the availability of this data. We use the lobbying reports provided as part of their bulk data facility. This is the full universe of available LDA-sourced reports (approximately 35,000 per year) that CRP has formatted, cleaned and modified. For example, the CRP reconciles different types of reports (taking account of amendments to original mid-year and end-of-year reports) and constructs lobbyist, firm and client identifiers. We conduct further cleaning and consolidation of the CRP identifiers in cases where the same individuals are split across different identifier codes. The LDA requires the reporting of lobbying spending above a $10,000 threshold and rounded to the nearest $20,000. In the case of self-filing organizations, the LDA requires the reporting of all expenses made on lobbying activities, including payments to outside entities as well as in-house employees.

The lobbying reports data is collapsed to the lobbyist-period level. Revenue is aggregated by lobbyist-period according to the 'unweighted' and 'weighted' measures defined in Section 2.3 of the main CHPATER. As a robustness check we re-estimated the main models and trimmed outliers at the 1 per cent and 2.5 per cent thresholds. This lead to only minor changes in coefficients and the results are available from the authors on request. Note that a manual check of outlier observations indicated that the majority of these high revenues belonged to well-known 'superstar' lobbyists.

Columbia Books Lobbyist Directory

Columbia Books publishes a comprehensive directory of Federal lobbyists under its suite of Lobbyists.info products. This directory initially began as a hard copy directory (titled

Washington Representatives) containing contact information on lobbyists and potential clients published in the late 1970s. Since this time Columbia Books has expanded the directory with further information on career histories, bio- graphical information, educational background and areas of expertise. The publisher then consolidated this directory in electronic form in 2006 as Lobbyists.info with daily updates and related supplementary databases.

This online version of the directory contains records on approximately 15,000 lobbyists. The career history information in Lobbyists.info includes the employer, job title and period of tenure for lobbyists' current and previous jobs, inclusive of private and public sector positions. We extract information on lobbyists who have had periods of political employment (that is, working as congressional staffers, in government agencies or as part of Presidential administrations) which is then matched by name into the CRP Lobbying Reports data.

LegiStorm Congressional Staffer Salaries

The second political employment database that we use is the LegiStorm Congressional Staffer Salaries (CSS) database. Based on Capitol Hill, LegiStorm is a company that aims to improve the availability of political information on the operations of the US Congress. For example, it provides easy-to-use versions of public data on Congressional remuneration; privately financed travel for Congress members and staff; financial disclosures; foreign gifts to members; and spending earmarks attached to bills.

The CSS database that we use is obtained by LegiStorm from published reports by the Secretary of the Senate and the Clerk of the House of Representatives. These reports are not actually made available in electronic form and LegiStorm takes the step of transferring the information from hard copy into an electronic format. As part of this process, LegiStorm also creates consistent set of identifiers for the staffers, offices and politicians that appear in the database.

LegiStorm's database contains information from late 2000 onwards. The main information provided is: staffer name; start and end dates for a given employment spell; office of employment within Congress; the job title or position; and the total salary amount for a given job spell. The full staffers database contains information on approximately 90,000 staffers over a nine-year period. This large number of staff is due to the inclusion of non-partisan institutional staff such as Capital Police. Our analysis focuses on the pool of approximately 40,000 staffers working in political or policy related offices over the 2000-2008 period.

Name Matching of Lobbyists

The full list of lobbyists represented in the CRP lobbying reports database is matched with individ- uals appearing in our two political employment databases, Lobbyists.info from Columbia Books and Congressional Staffer Salaries from LegiStorm. The name matching is implemented using a string- based algorithm which cleans the raw names for punctuation and shortened names (for example, ”JIM” becomes ”JAMES” and so on). The same algorithm is applied to each set of names and each politi- cal employment database is separately matched with the CRP lobbying reports list. The subsequent matches are then compiled into one list of political employees-turn-lobbyists. Middle names or initials are used as part of the name matching procedure where available.

We then score matches on their accuracy according to two criteria. Firstly, each match receives a 1-4 score based on how often a particular first or last name appears. We call this a ’name frequency’ score. Commonly occurring names such as ’SMITH’ are given a score of 4 while the least common names are given a score of 1. This process is repeated separately for first and last names to produce a 2-8 score. For example, a name such as ’JOHN SMITH’ is given an overall score of 8 since it is comprised of two common names while ’MILLICENT SMITH’ receives a score of 5 (+4 for the common last name but +1 for the relatively uncommon first name).

Secondly, we score the matches according to how well the timing of employment transitions links up across the data. Staffers leaving employment in the Congress should appear in the lobbying reports data within 1-2 years of their final employment spell. We construct a 0-1 flag for whether the timing of the transitions is consistent across the data. In the final stage of the matching we then manually check the accuracy of the matches characterised by very common names and/or inconsistent timing. We do this by manually cross-referencing names with online CVs and biographies. This final step of manual checks is done for all names with a name frequency score above 5. In order to evaluate the cut-off we ran regressions where we interacted the name score with our main variables of interest. These results indicated that measurement error in the name matching increased with the name score (ie: the Senate/House coefficient for the 8’s was lower than that for the 7’s which was in turn lower than the coefficient for the 6's). However, we did not pick up a significant difference in coefficients between the 5's and observations with a name score between 1-4.

In our final sample we also condition on the maximum number of politicians a lobbyist can be connected to, setting this to 5. Approximately 8% of lobbyists have experience working as personal staff for more than one politician. The number of lobbyists with experience working for 5 or more politicians represents less than 0.25% of this total. Furthermore, on inspection most lobbyists with 5 or more connections were classified officially in the Legistorm job titles as 'Shared Employees' rather than as exclusive members of personal staff.

Congressional Politicians

Our final major dataset contains information on the service and characteristics of politicians in the House and Congress since the beginning of the available lobbying reports in 1998. The specific data used is Stewart and Woon's (2009) compilation focusing on committee membership and updated periodically from the Congressional Record. This membership data here contains periods of service and reasons for exit (retirement, defeat for re-election etc) where applicable. The politicians appearing in the data have been allocated the ICSPR 'member id' that is common across political science studies in this area. We have name matched the list of politicians given in the political employment databases against the Stewart and Woon (2009) list using the same string-based algorithm developed for the lobbyist-level matching.

TABLE A2.1: DESCRIPTIVE STATISTICS – IN-HOUSE ORGANIZATIONS.

	In House
Panel A - Organizational Level	
Mean Number of Lobbyists	1.5
Total Revenue/Expenditures	786.1
Total Number of Firms/Organizations	3233
Panel B - Types of Lobbyists	
Revolving Door Lobbyist	0.185
Ex-Congressman	0.004
Ex-staffer:	0.137
- of politician serving pre-1998	0.017
- of politician serving post-1998	0.101
- of a congressional committee	0.019
Outside Congress	0.044
Panel C - Mean Revenue or Expenditure	
Weighted	
Revolving Door Lobbyists	292.5
-Ex-Congressmen	368.4
-Ex-staffers:	280.5
-Outside Congress	323.1
Other lobbyists	204.9
Unweighted	
Revolving Door Lobbyist	2319.8
-Ex-Congressmen	2134.0
-Ex-staffers:	2287.5
-Outside Congress	2438.3
Other lobbyists	1752.1
Panel D - Share of total industry expenditure, by type of lobbyist	
Revolving Door Lobbyist	0.222
-Ex-Congressmen	0.005
-Ex-staffers:	0.159
-Outside Congress	0.060
Number of Lobbyists (Total)	21,374
Number of Lobbyist-Period Observations	127,960

Note: Panel(A) based on 1998-2008 panel of organizations by period. Panels (B) and (C) based on 1998-2008 lobbyist-level panel. Length of each period is 6 months. Panel (C) presents annualised measures of expenditure per lobbyist. Panel (D) aggregates the weighted expenditures of lobbyists in order to calculate expenditure shares by type.

TABLE A2.2: AVERAGE EFFECTS OF REVOLVING DOOR CONNECTIONS ON LOBBYING REVENUE, WEIGHTED MEASURE OF REVENUE

	Dependent Variable: (log) revenue per lobbyist			
		Plus Party	Plus Chamber	Plus experience
	(1)	(2)	(3)	(4)
Number of Senators	0.20***	0.19***	0.20***	0.24***
	(0.07)	(0.07)	(0.07)	(0.07)
Number of Representatives	-0.01	-0.03	-0.03	-0.01
	(0.05)	(0.05)	(0.05)	(0.05)
Individual dummies	Yes	Yes	Yes	Yes
Time	Yes	No	No	No
Time*Party	No	Yes	No	No
Time*Party*Chamber	No	No	No	Yes
Lobbyist Experience	No	No	No	Yes
Individuals	1,113	1,113	1,113	1,113
Observations	10,418	10,418	10,418	10,418

Notes: This table presents the average effects of political connections on ex-staffers lobbying revenue. The dependent variable is the log of the **weighted revenue** generated from all the clients that an individual lobbyist serves in a time (semester) period. The two main independent variables are the number of senators and representatives that an individual lobbyist worked for previous to entering the lobbying industry *and are serving in Congress in that time period*. All regressions use a sample containing ex-staffers-turned-lobbyists and include both individual lobbyist dummies and time effects (i.e., semester dummies). Column 2 allows for different time effects for lobbyists connected to politicians in different parties (i.e., Democrats versus Republicans). Columns 3 and 4 allow for different time effects for lobbyists connected to politicians in different party/chamber combinations (i.e., Democrats in the Senate, etc.). Column 4 includes lobbyist experience (i.e., number of periods that a lobbyist appears in the sample) in quadratic form. Standard errors are clustered by lobbyist. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

TABLE A2.3: EFFECTS OF PARTY CONTROL.

Dependent Variable	(log) Weighted Revenue per Lobbyist	
	(1)	(2)
Number of Senators	0.25***	0.25***
	(0.07)	(0.07)
Number of Representatives	0.12**	0.11**
	(0.05)	(0.05)
House Dem Control*Dem	0.18***	0.19
	(0.06)	(0.06)
Senate Dem Control*Dem	0.14	
	(0.05)	
2001-2002*Dem		0.08
		(0.08)
2007-2008*Dem		0.18***
		(0.06)
Individual Dummies	Yes	Yes
Time	Yes	Yes
Lobbyist Experience	Yes	Yes
Individuals	1,113	1,113
Observations	10,418	10,418

Note: This table presents the effects of a party's control of Congress on the revenues of its affiliated lobbyists. One period of Democrat control in the House is considered (2007-2008) along with two periods of Democrat control in the Senate (2001-2002 and 2007-2008). In the second column we include the two periods of Democrat control in the Senate separately. We interact these party control dummies with whether the lobbyist is an ex-staffer for a Democratic politician. We also include the main variables in Table 2 as well as individual dummies, time dummies and lobbyist experience (and its square). Standard errors are clustered by lobbyist.

TABLE A2.4: ROBUSTNESS TESTS

	Party Control Dummies (1)	Time Trends (2)	Firm Dummies (3)
Number of Senators	0.26***	0.19*	0.28***
	(0.09)	(0.11)	(0.08)
Number of Representatives	0.11	0.02	0.04
	(0.08)	(0.08)	(0.06)
Individual dummies	Yes	Yes	Yes
Time*Party*Chamber	Yes	Yes	Yes
Lobbyist Experience	Yes	Yes	Yes
Ind.Dum*Party in Control	Yes	No	No
Lobbyist Time Trends	No	Yes	No
Firm Dummies	No	No	Yes
Individuals	1,113	1,113	1,113
Observations	10,418	10,418	10,418

Note: This table presents a number of robustness tests on Table 2. The main independent variables are as in Table 2.2. Every regression contains individual dummies, lobbyist experience in quadratic form, and different time effects for lobbyists connected to politicians in different party/chamber combinations. Column (1) includes the party control dummies used in Table A2.3. These dummies are interacted with the individual lobbyist dummies. Column (2) introduces 1,113 lobbyist-specific linear time trends. Column (3) includes 726 lobbying firm dummies. Standard errors are clustered by lobbyist.

TABLE A2.5: EFFECTS OF ENTRY INTO COMMITTEES.

	Dependent Variable: (log) revenue per lobbyist			
	(A) SENATE		(B) HOUSE	
	(1)	(2)	(3)	(4)
One Senator:				
in Finance	0.37***	0.32**		
	(0.13)	(0.14)		
in Appropriations	0.03	-0.13		
	(0.35)	(0.35)		
One Representative				
In Finance			0.37*	0.36*
			(0.20)	(0.20)
in Appropriations			-0.19	-0.22
			(0.22)	(0.22)
Dummies 2 Senators	Yes	Yes	No	No
Dummies 2 Representatives	No	No	Yes	Yes
Time*Party*Chamber	Yes	Yes	Yes	Yes
Lobbyist Experience	No	Yes	No	Yes
Individuals	534	534	644	644
Observations	4,457	4,457	4,727	4,727

Note: This table presents the effects of Congressmen joining important committees on the revenue of their connected lobbyists. The dependent variable is as in Table 2. The Senate regression sample contains only lobbyists connected to serving Senators. The displayed variable Finance takes value one when the connected and serving Senator has joined the Finance committee. We define similarly the variable Appropriations. The omitted group is being connected to one serving Senator in neither of these two Committees. For lobbyists connected to two serving Senators we also define and include dummies capturing whether one connected Senator joined the Finance or Appropriations committees. In practice, there are very few such cases and the estimated parameters are not displayed. The House sample and regressions are constructed equivalently. All regressions include individual lobbyists dummies and time effects (i.e. semester dummies). Columns (2) and (4) allow for different time effects for lobbyists connected to politicians in different party/chamber combinations (i.e. Democrats in the Senate, etc.) and also include lobbyist experience (i.e. number of periods that a lobbyist appears in the sample) in quadratic form. Standard errors are clustered by lobbyist.

CHAPTER 3: MINIMUM WAGES AND FIRM PROFITABILITY

Abstract

We study the impact of minimum wages on firm profitability, exploiting the changes induced by the introduction of a UK national minimum wage in 1999. We use pre-policy information on the distribution of wages to implement a difference in differences approach. Minimum wages raise wages, but also significantly reduce profitability (especially in industries with relatively high market power). This is consistent with a no behavioral response model where wage gains from minimum wages map directly into profit reductions. There is some weak suggestive evidence of longer-run adjustment through falls in net entry rates (JEL J23, L25).

JEL Classification Codes: J23, L25

Key Words: Firms, profits, wages.

3.1 Introduction

In debates on the economic impact of labour market regulation, much work has focused on minimum wages. Although the textbook competitive labour market model implies that wage floors raise the wages of the low paid and have a negative impact on employment (Borjas 2004; Brown 1999), the empirical literature is less clear-cut. Many studies have rigorously demonstrated that minimum wages significantly affect the structure of wages by increasing the relative wages of the low paid (e.g. DiNardo, Fortin and Lemieux, 1996).[32] However, in spite of the large number of studies, empirical evidence on employment effects is considerably more mixed (see the recent comprehensive review by Neumark and Wascher, 2007). Some have found the expected negative impact on employment[33], yet others have found no impact or sometimes even a positive effect of minimum wages on jobs.[34]

In the light of this, it is natural to ask how firms are able to sustain higher wage costs induced by the minimum wage. This paper explores the possibility that firm profit margins are reduced. A second possibility is that firms simply pass on higher wage costs to consumers in the form of price increases. However, there is scant evidence on this score.[35] Indeed, even with

[32] See also Lemieux (2006) for some recent evidence on the US and DiNardo and Lemieux (1997) for a comparison with Canada.

[33] See the discussion of time series studies in Brown, Gilroy and Kohen (1982) and Brown (1999) or the US cross-state panel evidence of Neumark and Wascher (1992) and the recent longer run analyses of David Neumark and Olena Nizalova (2007).

[34] Examples here are Dickens, Manning and Machin (1999) and Card and Krueger (1994).

[35] This was the conclusion of the survey on minimum wages and prices by Lemos (2008). For exceptions on restaurant prices see Aaronson (2001), Aaronson and French (2007) and Fougere, Gautier and le Bihan (2008). The only UK evidence to our knowledge is Wadsworth (2009) who finds limited effects on prices.

some positive price response, part of the higher wage costs may not be fully passed on to consumers and the minimum wage could eat directly into profit margins. A third possibility is that minimum wages may “shock” firms into reducing managerial slack and improving efficiency. We examine this productivity story but do not find any evidence for it.

Given this discussion, it is surprising that there is almost a complete absence of any study directly examining the impact of minimum wages on firm profitability. This is the focus of this paper. We adopt an identification strategy using variations in wages induced by the introduction of the national minimum wage (NMW) in the UK as a quasi-experiment to examine the impact of wage floors on firm profitability. The introduction occurred in 1999 after the election of the Labour government that ended seventeen years of Conservative administration. To date there is evidence that the NMW increased wages for the low paid, but had little impact on employment[36] and so this provides a ripe testing ground for looking at whether profitability changed.

Our work *does* uncover a significant negative association between the national minimum wage introduction and firm profitability. We report evidence showing wages were significantly raised, and firm profitability was significantly reduced by the minimum wage introduction. There is also some evidence of bigger falls in margins in industries with relatively high market power, but no significant effects on employment or productivity in any sector. Our findings can be interpreted as consistent with a simple no behavioral response model where wage gains from minimum wages map into profit reductions. There is a hint of a selection effect

[36] See Machin, Manning and Rahman (2003) and Stewart (2004).

in the longer-run as net entry rates fall in the most affected industries, but although the magnitude of the effect is nontrivial it is statistically insignificant.

The rest of the paper is structured as follows. In Section 3.2, we discuss a model of profit responsiveness to wage changes from which we derive our empirical strategy. Section 3.3 discusses the data and the characterisation of firms more likely to be affected by the minimum wage introduction. Section 3.4 gives the main results on wage and profitability effects and tests their robustness. Section 3.5 offers some further investigations using other datasets (care homes), other outcomes and sectoral heterogeneity. Section 3.6 concludes.

3.2 Motivation and Modelling Strategy

3.21. The Scope for Minimum Wages to Impact on Profitability

Following Ashenfelter and Smith (1979), consider a profit-maximizing firm employing a quantity of labor (L) at wage rate (W), using other factors at price R and selling its output at price P. Profits are maximized at $\Pi(\mathrm{W, R, P})$ given the values of W, R and P. The derivative of the profit function with respect to the wage rate is $\partial\Pi/\partial \mathrm{W} = -\mathrm{L}(\mathrm{W, R, P})$, the negative of the demand for labor. In turn, the second derivative is $\partial^2\Pi/\partial \mathrm{W}^2 = -\partial \mathrm{L}/\partial \mathrm{W}$.

In this setting, the introduction of a minimum wage (M) at a level above that of the prevailing wage reduces firm profits by $\Delta\Pi = \Pi(\mathrm{W, R, P}) - \Pi(\mathrm{M, R, P})$. Using a second-order Taylor series this can be approximated as:

$$\Delta\Pi \cong -\mathrm{L}\Delta \mathrm{W} + \frac{1}{2}\frac{\partial \mathrm{L}}{\partial \mathrm{W}}(\Delta \mathrm{W})^2$$

3.1)

where $\Delta W = M - W$. The terms on the right-hand side of equation (1) correspond to the "wage bill" ($-L\Delta W$) and "labor demand" ($\frac{1}{2}\frac{\partial L}{\partial W}(\Delta W)^2$) effects on profits. Note that equation (3.1) can be re-written as:

$$\Delta\Pi = -WL\left(\frac{\Delta W}{W} + \frac{\eta}{2}\left(\frac{\Delta W}{W}\right)^2\right) \qquad 3.2)$$

where $\eta = \frac{W}{L}\frac{\partial L}{\partial W} < 0$.

In a situation of "no behavioural response", that is no impact on labour demand, the second order effect in (2), ($\frac{\eta}{2}\left(\frac{\Delta W}{W}\right)^2$), is zero and the fall of profits that would result from the imposition of a minimum wage M is equal to the proportionate change in the wage multiplied by the wage bill. In the case of a labour demand effect the second term can offsets this profit loss to the extent that firms can substitute away from low-wage workers into other factors (e.g. capital).

Equation (3.2) also serves to illustrate the inverse relationship between a firm's initial wage and the post-policy change in its profits. It shows that, the lower the initial wage, then the greater the fall in profits associated with the imposition of a minimum wage. The difference-in-difference models we consider in our empirical modelling strategy (described below) will operationalize this idea by defining treatment groups of more affected firms, and comparison groups of less affected firms, based on their wages prior to the policy introduction.

Normalizing profits on sales revenues, S, to define a profit margin shows that, for the no behavioral response model, in a statistical regression context the coefficient on the increase in wages caused by the minimum wage $\left(\frac{\Delta W}{W}\right)$ should simply be equal to the share of the wage bill in total revenue ($\frac{WL}{S}$):

$$\Delta(\Pi/S) = -\theta\left(\frac{\Delta W}{W}\right) \qquad 3.3)$$

where $\theta = \frac{WL}{S}$.

More generally, to the extent there is substitution away from labor, the coefficient on the wage increase, θ, will be less (in absolute terms) than the (initial) wage bill share of revenue. Interestingly, we will show that our empirical results cannot generally reject the simple relationship in equation (3.3).

It is worth noting that this is consistent with the results in the rather different context of John Abowd's (1989) study of union wage increases and firm performance. Abowd estimates a version of equation (3.2) examining the effects of unanticipated increases in the wage bill ("union wealth") on the present discounted value of profits as reflected in changes in stock market values ("shareholder wealth"). He also finds that he cannot reject the simple model where the second order effect is zero. Abowd interprets this as evidence for strongly efficient union bargains as he focuses on a sample of unionized contracts. Strongly efficient (implicit) bargaining is also an alternative interpretation of our findings as well.[37]

---

[37] Although we find this explanation less plausible as the minimum wage mainly binds on those firms and sectors where unions are not present or, if they are, are very weak.

It is worth focusing on some of the economic issues underlying the adjustment mechanisms implicit in the second order term of equation (3.1). Obviously, the magnitude of these mechanisms depend on the elasticity of the labor demand curve, η. One element of this will be the degree to which labor is substitutable for other factors. Another will be the degree to which the higher wage costs can be passed on to consumers in the form of higher prices. For example, under perfect competition price equals marginal cost so all the wage costs are reflected in higher prices for consumers. In most oligopoly models, by contrast, mark-ups will fall as some of the wage increase is born by firms (see Appendix 3.A). Consequently, in our empirical work, we explicitly distinguish between industries with different degrees of product market competition as we expect heterogeneity in the minimum wage effects along this dimension (i.e. a larger effect in the less competitive industries).

The model focuses on the short-run responses when the number of firms is fixed, rather than in the long-run when the number of firms varies.[38] We believe that the short-run is still interesting as researchers cannot be sure how long is the long run (we look up to three years after the introduction of the minimum wage, so this is reasonably long run). Since firms that employ low-wage workers may well exit the market, so the relevant margin of adjustment will be more exit and less entry. We also examine this explicitly in our empirical analysis

Finally, when the product market is imperfectly competitive there may also be effects of the minimum wage on profitability in both the short-run and the long run. Appendix A in Draca et al (2008) discusses these models in some detail, but it is sufficient to note that

[38] Note that the short-run negative impact on profits will be larger in competitive labour markets than monopsonistic labour markets (see Card and Krueger, 1995). In the latter model, there is an offsetting positive effect on profitability when wages increase as worker turnover declines.

positive price cost margins are an equilibrium phenomenon in standard industrial organization models such as Cournot or differentiated product Bertrand. For example, consider a Cournot oligopoly where firms have heterogeneous marginal costs and constant returns to scale. Introducing a minimum wage has a differential impact on the firm employing more low skilled workers causing this firm to lose market share and suffer a fall in its price cost margin. However, so long as profits do not fall below the exit threshold it will remain in the market with lower profitability.

3.22 Modelling Strategy

The approach we take to identify minimum wage effects in the context of the above theoretical discussion is in line with the existing literature that analyzes the impact of national minimum wages. Typically, we look at a group of firms that were more affected by the NMW introduction than a comparison set of firms.[39] By "more affected", we mean where wages potentially rose by more due to the imposition of the minimum wage floor. This quasi-experimental setting enables us to compare what happened to profitability before and after NMW introduction in low wage firms as compared to what happened to profitability across the same period for a comparison group of firms whose wages were not affected as much (or at all) by the NMW introduction.

For ease of exposition, we begin our discussion of modelling by thinking in terms of a discrete treatment indicator of the minimum wage policy for a set of low wage firms with a pre-policy introduction wage, W^{pre}, beneath the minimum wage threshold M. A treatment

---

[39] See, amongst others, Card's (1992) analysis of state variations in low pay incidence to identify the employment impact of the US federal minimum wage, or Stewart's (2002) similar analysis of regional variations in the UK NMW.

indicator variable can be defined as T = 1 for below minimum wage firms (where $W^{pre} < M$) and T = 0 for a set of firms whose pre-policy wage exceeds the threshold.[40]

We can evaluate the impact of minimum wages on firm profitability by comparing what happens before and after minimum wage introduction across these treatment and control firms. For this procedure to be valid, we first need to establish that our choice of affected firms behave as we would expect in response to NMW introduction. The expected response would be that wages rise by more in the T = 1 firms before and after introduction as compared to the T = 0 firms.

A difference-in-difference estimate of the wage impact of the NMW is $(\overline{w}^{T=1}_{NMW=1} - \overline{w}^{T=1}_{NMW=0}) - (\overline{w}^{T=0}_{NMW=1} - \overline{w}^{T=0}_{NMW=0})$, where w = ln(W), NMW is a dummy variable equal to 1 for time periods when the NMW was in place (and 0 for pre-policy periods) and a bar denotes a mean. For example, $\overline{w}^{T=1}_{NMW=1}$ is the mean ln(wage) for the treatment group in the post-policy period. This difference-in-difference estimate is just the simple difference in means unconditional on other characteristics of firms. It can easily be placed into a regression context. If T = 1 for firms with a pre-policy ln(wage), $w_{i,t-1}$, less than the ln(minimum wage), mw_t, and 0 otherwise, we can enter the indicator function $I(w_{i,t-1} < mw_t)$ into a ln(wage) equation for firm i in year t as follows:

$$w_{it} = \alpha_1 + \beta_1 X_{it} + \delta_1 Y_t + \theta_1 I(w_{i,t-1} < mw_t) + \psi_1 [I(w_{i,t-1} < mw_t) * NMW_t] + \varepsilon_{1it} \tag{3.4}$$

3.4)

[40] We also consider various continuous measures of treatment intensity discussed below.

where X is a set of control variables, Y denotes a set of year effects (hence a linear term in NMW_t does not enter the equation since it is absorbed into the time dummies) and ε_{1it} is a random error. Here the regression corrected difference-in-difference estimate of the impact of NMW introduction on the ln(wage) is the estimated coefficient on the low wage treatment dummy in the periods when the NMW was in operation, ψ_1.

After ascertaining whether the NMW impacts on wages in the expected manner we move on to consider whether profitability was affected differentially between the treatment group firms (T = 1) and comparison group firms (T = 0). We look at unconditional and conditional difference-in-difference estimates in an analogous way to the wage effects. Thus, we can estimate the unconditional difference-in-difference in profit margins, defined as the ratio of profits to sales Π/S, as $\left[\left(\frac{\bar{\Pi}}{S}\right)^{T=1}_{NMW=1} - \left(\frac{\bar{\Pi}}{S}\right)^{T=1}_{NMW=0}\right] - \left[\left(\frac{\bar{\Pi}}{S}\right)^{T=0}_{NMW=1} - \left(\frac{\bar{\Pi}}{S}\right)^{T=0}_{NMW=0}\right]$ and the conditional difference-in-difference, ψ_2, from the regression model:

$$\left(\frac{\Pi}{S}\right)_{it} = \alpha_2 + \beta_2 Z_{it} + \delta_2 Y_t + \theta_2 I(w_{i,t-1} < mw_t) + \psi_2 [I(w_{i,t-1} < mw_t) * NMW_t] + \varepsilon_{2it} \quad 3.5)$$

where the controls are now Z and ε_{2it} is the error term.

If we compare the econometric models (4) and (5) to the economic models of (1) through (3), we see immediately that the no behavioral response model corresponds to a restriction on the coefficients in equations (4) and (5), i.e.

$$\psi_2 = -\theta\psi_1 \quad 3.6)$$

We present formal tests of this restriction in the empirical section.

The main issue that arises with any non-experimental evaluation of treatment effects is, of course, whether the comparison group constitutes a valid counterfactual. The key conditions are that there are common trends and stable composition of the two groups (see Blundell, Costa-Dias, Meghir and Van Reenen, 2004). Much of our robustness analysis below focuses on whether these two conditions are met: for example, by examining pre-policy trends and carrying out pseudo-experiments (or falsification tests) in the pre-policy period.

3.3 Data

3.31 Basic Description of FAME Data

Accounting regulations in the UK require private firms (i.e. those unlisted on the stock market) to publicly report significantly more accounting information than their US counterparts. For example, even publicly quoted firms in the US do not have to give total employment and wage bills whereas this is required in the UK.[41] Accounting information on UK companies is stored centrally in Companies House. It is organised into electronic databases and sold commercially by private sector data providers such as Bureau Van Dijk from whom we obtained the FAME (Financial Analysis Made Easy) database.[42]

The great advantage of this data is that is covers a much wider range of companies than is standard in firm level analyses and, in particular, it includes firms who are not listed on the stock market. This means we are able to include many of the smaller and medium sized firms that may be disproportionately affected by the NMW. Furthermore, the data also covers non-

[41] The lack of publicly available information on private sector firms and on average remuneration may be a reason for the absence of US studies in this area.

[42] FAME is the UK part of BVD's AMADEUS dataset of European company accounts used by many authors (e.g. Bloom and Van Reenen, 2007).

manufacturing firms where many low wage workers are employed. By contrast, plant level databases in the UK and US typically cover only the manufacturing sector[43] and do not have as clear a measure of profitability as exists in the (audited) company accounts. However, UK accounting regulations do have reporting exemptions for some variables for the smaller firms so our analysis is confined to a sub-sample who do report the required information.[44]

Since FAME contains annual accounting information, we have firms reporting accounts with different year-end dates. Since the NMW was introduced on April 1st 1999, we therefore consider the sub-set of firms who report their end of year accounts on March 31st of each year (these are firms who report in the UK financial year). The accounting period for these firms will match exactly the period for which the NMW was in force. Around twenty-one percent of firms in FAME who have the accounting data we require report on this day, which corresponds to the end of the tax year in the UK.[45]

We use data on profits before interest, tax and depreciation from the FAME database and model profitability in terms of the profit to sales ratio. There is a long tradition in firm-level profitability studies to use this measure, as it is probably the best approximation available in

[43] The Annual Business Inquiry (ABI) database does cover non-production sectors, but this database is not available until the late 1990s. The US Longitudinal Research Database (LRD) only covers manufacturing.

[44] These firms will tend to be larger than average as the very smallest firms have the least stringent reporting requirements.

[45] If we estimated our basic models on the whole FAME sample irrespective of reporting month we obtained very much the same pattern of results as our basic findings in Table 3.2 below. The estimated effects were a little smaller in magnitude, most likely because of attenuation towards zero owing to measurement error in defining treatment.

firm-level accounts data to price-cost margins.[46] To allow for capital intensity differences we also control for firm-specific capital to sales ratio.[47]

3.32 Other Data

We have also matched in industry-level variables aggregated up from the Labour Force Survey (similar to the US CPS). These are used as control variables in the analysis and include (at the three-digit industry level) the proportion of (a) part-time workers, (b) female workers and (c) union members. We also include skills proxied by the proportion of all workers who have college degrees in a particular region by two-digit industry cell. The control variables in the regression models also include a set of region, one-digit industry and time dummies. Exact variable definitions are given in the Data Appendix. Appendix Table 3.B1 shows the characteristics of the treatment and comparison groups for each model.[48]

Finally, the magnitude of the minimum wage increases over our "Policy On" period should be clarified. This period lasts from April 1st 1999 until March 31st 2002 (the end of our sample). Along with the introduction of the minimum wage, there were two upratings of the

---

[46] For example, see Machin and Van Reenen (1993) and Slade (2004). Although there are many reasons why accounting and economic profits may diverge (Fisher and McGowan, 1983), there is much evidence that they are on average highly positively correlated. The relationship between the profit-sales ratio and price-cost margins will also break down if there are not constant returns to scale. In this case, controlling for capital intensity is important in allowing for differential fixed costs across firms and that is what we do empirically in the regression-corrected difference in difference estimates.

[47] We also checked that dropping the capital sales ratio did not change the results as some of the effect of the NMW may have come from firms substituting away from more expensive labour towards capital equipment.

[48] Interestingly the profitability of low wage firms is higher at the median and mean than comparison group firms. This is not true for firms as a whole where there is a positive correlation between average firm wages and profits per worker (e.g. Van Reenen, 1996). It is because we are focusing on the lower part of the wage distribution that this correlation breaks down.

minimum during this time. The first occurred in October 2000 and saw the minimum wage rise by 10p to £3.70. The second uprating a year later was more substantial taking the minimum up to £4.10. Together these upratings constitute a 13.9% increase in the minimum between 1999 and 2002.[49] Small cell sizes prevent us from estimating separate models for the 2000 and 2001 upratings.[50]

3.4 Defining Treatment and Comparison Groups

3.41 Basic Approach

FAME has a total remuneration figure that can be divided by the total number of employees to calculate an average wage.[51] This creates a challenge in terms of defining our treatment and comparison groups since any given level of average wages is, in principle, compatible with a range of different within-firm wage distributions. This makes it hard to measure accurately how exposed each firm's cost structures are to the wage shock brought about by the minimum wage. For example, any continuous measure of treatment intensity based on the firm average wage is inevitably coarse.

We have used information from FAME, the Labour Force Survey (LFS) and the Workplace Employment Relations (WERS) to both construct and validate our treatment group indicators. Specifically, the main results use average firm wages from FAME to define our

[49] By contrast, the consumer price index grew by 6.3% over the same period.

[50] For example, less than 9% of firms report annually on September 30th (i.e. the 12 months immediately before the October upratings).

[51] In almost all firms in the data we use, employment refers to average employment over the accounting period. Firms can report employment at the accounting year or the average over the year, but the overwhelming number of our firms report averaged employment.

treatment and comparison groups, but we also use LFS information for the industry level analysis of entry and exit. We use within-establishment information from matched worker-establishment data in WERS to consider the association between low pay incidence and average wages to assess the effectiveness of this empirical strategy.[52]

To investigate the impact of the minimum wage we have defined our treatment group, T, based upon average remuneration information from FAME. For our main initial analysis we define T = 1 for firms with average remuneration of less than £12,000 in the accounting year prior to minimum wage introduction ("low wage firm").[53] Average remuneration in the treatment group for this threshold is £8,400 which, after allowing for a deduction for non-wage costs (such as employers' payroll tax, pension contributions, etc), is equivalent to a £3.90 hourly wage for a full-time worker and is close to the NMW (introduced at £3.60 per hour). For our research purposes, the key issue is that the wages of firms beneath the threshold we choose have a significant wage boost from the NMW relative to higher wage firms and we consider this in detail in our analysis. One aspect of this is that we have extensively experimented with the threshold cut-off and we discuss this in detail below. In the analysis presented below, we also look at associations with the pre-policy average wage in the firm. This gives a continuous

[52] Unfortunately, direct linking of data of WERS and FAME is not possible due to confidentiality restrictions.

[53] In earlier versions of this paper we also combined the low wage firm information with industry-region "cell" data on the proportion of workers beneath the minimum wage in the year before it came in. Using LFS data, we defined a low wage industry-region cell if more than 10% of workers in the given firm's two-digit industry by region cell in the pre-policy period are paid below the minimum wage. In practice this made little difference to the overall pattern of results and so we do not report this material (see Draca et al, 2008, for all the results).

indicator that we can use to compare with the binary treatment variables based upon being beneath a particular wage threshold.

3.42 The Usefulness of Average Wages to Define Treatment

How accurate are these treatment group definitions at identifying firms most affected firms by the minimum wage regulation? This hinges on how segregated low-wage workers are between firms. Our threshold-based definition will be more effective if sub-minimum wage employees are concentrated in particular firms at the lower end of the wage distribution.

To assess the usefulness of the approach we adopt we look at segregation and wages in the 1998 cross-section of the British Workplace Employment Relations Survey (WERS)[54]. This contains matched worker and establishment data that allows us to look at within-workplace wage distributions and explore the association between average wages and the intensity of low-wage workers. For 26,509 workers in 1,782 WERS workplaces we computed the proportion of workers paid less than £3.60 per hour (the value of the minimum wage when introduced in 1999) and the average hourly wage in the workplace. There is a strong, negative association between the two variables (a correlation coefficient of -0.61, p-value < 0.001). In Figure 3.1 we plot the proportion of workers paid at or below the minimum wage against the establishment's average annual wage. This proportion of minimum wage workers tapers off rapidly after an average annual wage of £10,000, supporting the idea that exposure to the minimum wage can be proxied by using an average wage threshold that is around this level. Workplaces with average annual wages of £12,000 or less (our main threshold defining the treatment group) contain 87%

[54] WERS is a stratified random sample of British establishments and has been conducted in several waves since 1980. It has been extensively used by economists and industrial relations experts to study a range of issues. Culley, Woodland, O'Reilly, Dix and et al. (1999) give details of the survey

of all minimum wage workers. These patterns give some support to our idea that "at risk" group of minimum wage workers are concentrated in firms that pay low average wages.

3.5 Main Results

3.51 Changes in Wages Before and After the Introduction of the National Minimum Wage

It is important to see whether we are able to observe a clear change or "twist" in the firm average wage distribution as the minimum wage was introduced. To consider this, we started our analysis by calculating the change in average wages in the year immediately before and immediately after NMW introduction for every firm at each percentile of the pre-policy firm wage distribution. If the firms in the FAME data exhibit some of the low pay patterns outlined above for WERS, the minimum wage introduction should raise average firm wages by more in low wage firms. Thus, we would expect there to be larger changes in firm wages for the lowest percentiles of the distribution.

The results given in Figure 3.2 very clearly confirm this hypothesis. In the post-NMW introduction year from April 1 1999 to March 31 2000 (labelled "1999-2000 Change", and denoted by the solid line), the wage change tapers off steadily beyond the lowest decile of the firm average wage distribution. After the 13th percentile, firms appear to have had a similar increase in nominal wages of around 5.6%. Importantly, there is no evidence of much faster wage growth for the bottom decile in the pre-policy year (labelled "1998-1999 Change", and denoted by the dotted line). In fact, wage growth in the bottom thirteen percentiles was on average 2.6% in the 1998-1999 financial year compared to 9.9% in the following year. A spike is seen for the bottom few percentiles of the wage distribution in both years, which is consistent with the notion of some transitory measurement error at the low end of the wage distribution generating mean reversion in both periods. Reassuringly, the general picture follows a similar pattern to that found for individual-level wage data (Dickens and Manning, 2004) and again provides encouraging evidence that our definition of the treatment group is useful.

It is critical that we identify wage effects from the treatment group definitions, so that our analysis of profitability consequences is validated by the minimum wage introduction having a bigger 'bite' on low wage firms. To make this a tighter definition we have also defined the comparison group to be those firms with average wages above the £12,000 treatment threshold, but less than £20,000 (the median firm wage) by removing any firms with above £20,000 average wages from the main analysis. We do so since these firms are quite different in terms of their characteristics and therefore subject to different unobservable trends from the treatment group. We are careful to test for the sensitivity of the results to definitions of these thresholds.

3.52 Firm-Level Estimates: Wages and Profitability

The upper panel of Table 3.1 presents unconditional difference-in-differences in the mean ln(wage) for the discrete categorization of treatment and comparison groups, for the three years before and after NMW introduction.[55] It is evident that wages rose significantly faster amongst the low wage firms when the minimum wage became operational. Wage growth across the pre- and post-NMW three year time period was higher at 22.9 log points in the low initial wage group (T = 1) as compared to wage growth of 11.8 log points in the higher initial wage group (T = 0). The difference-in-difference of 11 percent is strongly significant in

[55] Note that we are looking across the six financial years from April 1 1996 to March 31 2002 (three years before the policy and three years afterwards). In Figure 2, we simply looked one year before and after the policy introduction.

statistical terms. This is consistent with the hypothesis that the NMW significantly increased wages for low wage firms[56].

An analogous set of descriptive results is presented for firm profitability in Panel B of Table 3.1. It is clear that, whilst profit margins fell by 0.039 between the pre- and post-NMW periods in the pre-NMW low wage firms, they only fell by 0.012 in the pre-NMW higher wage firms. Thus, there is a negative difference-in-difference of -0.027. This difference is statistically significant and is preliminary evidence that profit margins were squeezed in firms that were "at risk" from the introduction of the minimum wage.

Comparing these results with the simple models in Section 3.2, we find that no behavioral response model does surprisingly well. Using the average wage bill to sales ratio of 0.27 (see Table 3.B1), the implied change of profit margins using the estimated wage gains in Table 1 and equation (3.3) is -0.030 (= -0.111*0.27). This is only slightly above the empirically estimated profitability reduction of -0.027 in Table 3.1, suggesting only minor offsetting adjustments (the second-order term in equation (3.2)). Below, we will see that this conclusion broadly holds up to more rigorous econometric testing.

Table 3.2 reports results from statistical difference-in-difference wage and profitability regressions that additionally control for firm and industry characteristics. The upper panel (A) of the Table shows results for the binary low wage firm indicator, whilst the lower panel (B) uses a continuous measure, the negative of the pre-policy average wage (reporting the negative so as to

[56] As we saw in Figure 3.1, in 1998 (the year prior to the introduction of the National Minimum Wage in 1999), on average 25% of workers in the treatment group were at or below the minimum wage compared to 3% in the comparison group. Based upon this 22 percentage point difference, our coefficients would have to be scaled up by a factor of 4.5 if we considered the more radical experiment of switching a firm from having *none* of its workers covered to having *all* of its workers covered by the minimum wage.

have signs on coefficients that are consistently defined with the low wage dummy). The basic pattern of results from the unconditional models of Table 3.1 are confirmed in these conditional specifications. For the binary indicator in the upper panel, the estimated effects show a 9.4 percentage point increase in wages and a 2.9 percentage point decrease in profitability (similar to Table 3.1). The same pattern of results is observed for the (negative of the) continuous pre-NMW wage, reported in panel B. There is a significant positive connection between wage growth and the negative of the pre-NMW wage and a significant negative association with profitability. When compared to average profits in the low-wage firms in the pre-policy period, the results for the binary low-wage firm model imply a sizable 22.7 per cent (-0.029/0.128) fall in profit margins. The P-values from F-tests of the no behavioral response model are at the bottom of each panel and again indicate that we cannot reject the simple model underlying equation (3).

3.53 Further Probing of the Baseline Results

There are many reasons to probe these baseline results more deeply. The first, and obvious, reason is to judge the sensitivity of our definition of pre-policy low wages. Because we do not have data on the individual workers within our FAME firms, we rely on pre-policy low wage status as being a function of the average wage in the firm. This is less than ideal, even though we have (at least partially) validated its use above with the WERS data, and it is important to study whether the results are robust to alternative ways of defining the threshold between treatment and comparison groups.

We therefore re-estimated the models in Table 3.2 for a range of different wage thresholds, running from an average wage of £10,000 at £1,000 intervals up to £15,000. The results are reassuring in that they all establish a significant NMW effect of reducing profit margins, with magnitude of the impact varying and becoming slightly larger (in absolute terms) for lower thresholds as we would expect (so there is a bigger impact on the very low wage firms).[57]

A second possible concern is that our results are simply picking up a relationship between changes in profit margins and initial low wage status that exists, but has nothing to do with the NMW introduction. We have thus looked at estimates, structured in the same way, from periods *before* the NMW was introduced. One such 'placebo experiment' is reported in Table 3.3 where we examine an imaginary introduction of the NMW on April 1st 1996 (instead of April 1999) and repeat our analysis of wage and profitability changes. Table 3 very much reinforces the results reported to date, as we are unable to find any difference in margins between low and high wage firms in the period when the policy was not in place. This is consistent with the NMW introduction being the factor that caused margins to fall in low wage firms.

A related issue is the possibility of pre-sample trends (possibly due to mean reversion) in the wage model. If initially low wage firms had lower than average profitability growth even in the absence of the policy this would be conflated with the causal effect of the NMW impact on profits. The evidence from Table 3.3 suggested that there is no trend for wages or profitability in the pre-policy period. Nevertheless, we investigated this issue in more detail by estimating the profits model of Table 3.2 with a rolling threshold from £10,000 to £15,000 for both the policy and pseudo-experiment periods. That is, we estimate the model for thresholds at

[57] The profitability impacts for the different T = 1 thresholds were: -0.029 (0.014) for £10,000; -0.027 (0.013) for £11,000 ; -0.029 (0.012) for £12,000 ; -0.024 (0.010) for £13,000; and -0.014 (0.009) for £14,000.

each £100 interval in this range and plot the coefficients (see Figure 3.3). In the policy-on period there is a consistently negative effect of around 2-3% no matter how we draw the exact profit threshold. By contrast, in the pre-policy period there is essentially a zero effect with the point estimates actually positive and around 1%.

Draca et al (2008) report a number of further robustness tests. First, a statistical matching technique by trimming the sample according to the propensity scores of the treatment and comparison groups did not affect the pattern of results.[58] As discussed earlier our sample seems well chosen with relatively few observations needing to be trimmed to ensure common support. More importantly, the estimated effect of the policy on wages and profitability are significant and similar to those in the baseline Low Wage Firm specification.[59] Second, we included a full set of three-digit industry time trends. Although this is a strong test, the profitability effect was almost identical when these industry time trends were included with an estimate of -0.032 (0.015).

[58] The basic method used is that of Heckman, Ichimura and Todd (1997) where propensity scores are estimated and the sample then trimmed to exclude poorly matched observations without common support. To generate the propensity scores, we used a probit model that included all the control variables used in Table 3.2. We trimmed at the 1st percentile of the treatment group and the 99th percentile.

[59] Few observations are lost under propensity score matching because the comparison group is already chosen to be of relatively low wage firms (under £20,000 average annual wages). If we had used the entire FAME sample (including firms with average wages of over £20,000) we would have had to lose the vast majority of the sample to ensure that the comparison group had common support with the treatment group. Results are not presented for the pre-policy average wage since that is a continuous variable. If, however, the specification including that variable was estimated on the trimmed sample from columns (2) or (3) this produced very similar results to the baseline estimates of Table 3.2.

3.6 Further Investigation of the Minimum Wage Effect

The baseline results of Section 3.4 show very clearly that low wage firms in the FAME data experienced faster wage growth coupled with falling profit margins before and after the introduction of the UK National Minimum Wage. The results also seem consistent with the no behavioral response theoretical model introduced in Section 3.2 above. The model has a number of other salient features that we explore more fully in this Section, in an attempt to understand the effect of minimum wages on firm profitability and mechanisms that underpin the negative effect our baseline results have uncovered.

3.61. Minimum Wages and Profitability in UK Residential Care Homes

In this sub-section we look at the wage and profitability effects of the minimum wage in a rather different context, UK residential care homes.[60] There are three reasons to focus on care homes to juxtapose with the FAME results. First, it is a very low wage sector so offers a good testing ground for studying minimum wage effects on profitability and other economic outcomes.[61] Second, the sector is price regulated so one of the margins of adjustment (passing on higher wage costs in higher prices) is constrained. Finally, we have individual level data so can observe the entire within-firm wage distribution in this exercise, something we could not do in the FAME dataset.

---

[60] To date these data have mostly been used for studies of minimum wage effects on wages and jobs (e.g. Machin, Manning and Rahman, 2003), but see also Machin and Manning's (2004) test of competitive labour market theory.

[61] Prior to the minimum wage introduction in April 1999 average hourly wages were very low in the sector (at around £4 per hour). On average, 32.2% of workers were paid below the incoming minimum wage with this figure falling to 0.4% after the introduction of the policy.

The more sophisticated definition of treatment we are able to use is the initial firm wage gap relative to the minimum, namely the proportional increase in a firm's wage bill required to bring all of its workers up to the minimum wage. This variable, GAP, is defined as:

$$GAP_i = \frac{\sum_j h_{ji} \max(W_{ji}^{min} - W_{ji})}{\sum_j h_{ji} W_{ji}} \qquad 3.7)$$

where h_{ji} is the weekly hours worked by worker j in firm i, W_{ji} is the hourly wage of worker j in firm i, and W_{ji}^{min} is the minimum wage relevant for worker j in firm i .

For care homes, we do not have accounting data and so the profit variable we study is a derived one based on total revenues less total costs. Total revenue of each home is measured directly as the product of the number of beds, the home-specific average price of beds and the home occupancy rate. Total costs are calculated by dividing the total firm wage bill by the share of labour in total costs.[62] Home profitability is then defined as the ratio of profits to revenue.

We therefore estimate the following care homes specification:

$$\Delta\left(\frac{\Pi}{S}\right)_{it} = \eta_0 + \eta_1 GAP_{i,t-1} + \eta_2 Z_{i,t-1} + \xi_{it} \qquad 3.8)$$

where ξ_{it} is the equation error. Under the no behavioral response model the coefficient on GAP (η_1) should be equal to the wage bill share of revenues.

---

[62] Total sales and profits are not reported directly in the care homes data. We calculated them from the underlying home-specific components. Sales (*S*) is calculated as Occupancy Proportion* Number of Beds * Average Price (all reported in the survey). The wage bill (*WB*) and the share of labor in total costs (*SHARE*) are also reported directly in the data. We can then calculate total costs (*TC*) as the ratio of the wage bill to the labor share (*WB/SHARE*). Profits are then simply sales less total costs (*S* - *TC*). Profitability is the ratio of profits to sales, (*S* - *TC*)/*S*.

Table 3.4 presents estimates of home-level wage change and profitability change equations for the period surrounding NMW introduction (1998-99). The upper panel (A) focuses on wages, and presents results showing that wages clearly rose by more in homes with a larger pre-NMW wage gap. The lower panel (B) shows profitability estimates, where the coefficient on the pre-NMW wage gap variable is estimated to be negative and significant. In the column (2) specification with controls the coefficient is -0.492. Thus there is clear evidence of profitability falls in homes that were more affected by the minimum wage introduction. This very much corroborates the FAME findings of the previous section.

There was also some evidence that wages rose more in the pre-policy period (1992-93) in homes with a bigger initial wage gap.[63] Nevertheless, the relationship is much weaker in the earlier period so the trend-adjusted estimate is statistically significant and large in magnitude (at 0.678). Under the no behavioural response model, the coefficient on the initial wage gap measure should equal the share of the wage bill in sales. The (trend adjusted) point estimate on the wage gap term in the profitability equation turns out to be -0.396 for the model with controls (and -0.343 for the no controls specification), which in absolute terms is very close to the wage bill to sales ratio in our sample of care homes (0.398). Hence, like the FAME results the magnitude of the estimated impact in care homes is very much in line with what we would expect from the simple no behavioral response model.

[63] We define a counterfactual minimum wage at the same percentile of the wage distribution as the real 1999 minimum, so we can compute a *GAP* measure for the earlier pre-policy time period. Note that this is the only previous wage change information that exists, as the data was not collected in other (non-election) years.

3.62 Sectoral Heterogeneity: Industries With High and Low Market Power

As noted in Section 3.2, a condition for the existence of long-run effects of minimum wages on profitability is that there is some degree of imperfect competition in the product market. To examine this idea in Table 5 we split industries into "high" and "low" competition industries based on a proxy for the Lerner Index (constructed as in Aghion et al, 2005). Consistent with the idea of imperfect competition, the effects of the NMW policy on profitability were stronger in the less competitive sectors (defined as those with above the median value of three-digit industry Lerner index). Table 3.5 shows that the impact of the policy on wages was not so different (10.9% vs. 8.1%). By contrast, the effect of the minimum wage on profitability was almost two and a half times as large in the less competitive industries as in the more competitive sectors (as well as being significant only in the more competitive sectors).

Under perfect competition, an industry facing a common increase in marginal costs will pass on the higher wage costs in the form of higher prices to consumers. In less competitive sectors, however, firms will generally adjust by reducing their profit margins, rather than just through prices. Therefore, the evidence in Table 3.5 is consistent with the idea that the strongest effects of the NMW on profitability will be in the less competitive sectors.

3.63 Effects of Minimum Wage on Other Outcomes: Employment, Productivity, Exit and Entry

We also examined the effect of the NMW policy on other firm outcomes in the lower part of Table 3.5, again split by high and low market power sectors. We do not find any significant negative effects on employment, consistent with some of the minimum wage literature (e.g. Card and Krueger, 1994). The presence of no significant employment effect is also consistent with our tests of the no behavioral response model. Similarly, there does not appear to be any effect of the policy introduction on labor productivity (as predicted by the "shock" theory).

The FAME database identifies four categories of inactive firms, namely firms that are dissolved, liquidated, in receivership or currently non-trading.[64] Hence, we have defined all firms in these categories as “exiting” firms. We examine three year death rates for a cohort alive April 1st 1999 (i.e. did they exit by March 31st 2002) compared to a cohort alive on April 1st 1996 (i.e. did they exit by 1999). In the final row of Table 3.5 there is no evidence of any faster increase in exit rates in initially low wage firms following the minimum wage introduction either in the whole sample or in sub-sectors. The same is true in models of the probability of closure of care homes (see Machin and Wilson, 2004).

There are two possible problems with this firm-level analysis of exit. First, we ignore the possible entry-deterring effect of the minimum wage may, and second, there may be pre-policy trends.[65] Table 6 takes both of these into account. Obviously, we cannot implement this at the firm level, as entrants do not have a pre-policy wage for the entrants. However, we can examine an alternative dataset containing all entrants and exits in each three-digit sector (from the Department of Trade and Industry’s VAT Registration Database).[66]

---

[64] So exits by takeover are *not* coded to be unity in this definition as takeovers may be regarded as a sign of success rather than failure. Re-defining the dependent variable to be unity if the exit is to a takeover does not change the qualitative nature of the results.

[65] Running the pseudo-policy experiment of Table 3.3 gave a coefficient on the policy variable of 0.021 with a standard error of 0.106 for employment and 0.077 with a standard error of (0.053) for productivity.

[66] Unlike the firm data, we cannot distinguish between exit due to takeover and exit due to bankruptcy. Appendix table 3.B2 describes some key features of these data.

The three Panels of Table 3.6 show one-year entry rates, one-year exit rates and the difference between the two ("net entry") three-digit industries. Column (1) shows estimated coefficients on a pre-NMW low pay proportion in the period surrounding NMW introduction. Column (2) does the equivalent experiment for an imaginary/placebo policy (as in Table 3) introduced in 1996 and column (3) presents the trend-adjusted difference in differences. Although the first row shows that entry rates appear to perversely increase for low wage firms after the minimum wage, there does appear to be some positive pre-policy trend in column (2) suggesting a negative trend-adjusted effect of the NMW policy on entry. Similarly, trend adjusted exit rates in Panel B are 1.5 percentage points higher after the minimum wage was introduced. The final row shows that trend-adjusted net entry rates had fallen by about 5.1 percentage points in the low wage industries after the NMW introduction This effect is large in magnitude, but not statistically significant. These results do hint that in the long-run a margin of adjustment may be in the dimension of lower rates of net entry into the sectors most affected by the NMW.[67] There is little within firm change, but the margin of adjustment may be through the long-run number of firms.

3.7 Conclusions

This paper considers a very under-studied research question on the economic impact of minimum wages by looking at empirical connections between minimum wage legislation and firm profitability. Using the quasi-experiment of the introduction of a national minimum wage to the UK labour market in 1999, we utilise pre-policy information on the distribution of wages to

[67] Our further indications indicated that there were minimal differences in entry and exit rates between high and low market power industries. For example, when split by market power the corresponding estimates for column (1), Panel (A) in Table 6 were 0.025 (0.022) for high and 0.019 (0.020) for low.

construct treatment and comparison groups and implement a difference in differences approach. We report evidence showing wages were significantly raised, and firm profitability was significantly reduced by the minimum wage introduction. There is also some evidence of bigger falls in margins in industries with relatively high market power, but no effects on firm employment or productivity. Somewhat surprisingly, our findings are consistent with a simple "no behavioral response" model where wage gains from minimum wages map into profit reductions. There is a hint that the long-run adjustment may be through lower rates of net entry.

There are, of course, a number of caveats to our results. It would have been useful to have data on prices and quality to see if these may also have adjusted in response to minimum wages.[68] It would also be useful to have more information on the within firm distribution of workers in other sectors besides care homes. A fuller integration of theory and empirical work in the context of imperfect competition in both product and labor markets is another fruitful research area for the future. Overall, though given the total sparsity of evidence of the impact of minimum wage floors on firm profitability, we believe this study is an important contribution looking at the impact of labor market regulation on *firms* as well as the more developed and extensive evidence base that exists studying the impact on individuals.

[68] Although there is no evidence for these effects in the care homes sector, as it is heavily regulated (see Machin, Manning and Rahman, 2003).

REFERENCE LIST

Aaronson, D (2001) 'Price Pass Through and the Minimum Wage', Review of Economics and Statistics, 83, 158-169.

Aaronson, D and French, E (2007) "Product Market Evidence on the Employment Effects of the Minimum Wage" Journal of Labor Economics, 25 (1) 167-200

Aghion, P; Bloom,N; Blundell, R; Griffith,R and Howitt, P (2005), "Competition and Innovation: An inverted U relationship", Quarterly Journal of Economics,120, 701-728.

Ashenfelter, O and Smith,R (1979) 'Compliance with the Minimum Wage Law', Journal of Political Economy, 87, 333-50.

Besley, T and Burgess, R (2004) 'Can Labor Regulation Hinder Economic Performance? Evidence from India', Quarterly Journal of Economics, 119, 91-134.

Bloom, N and Van Reenen, J (2007) 'Measuring and Explaining Management practices across firms and nations', Quarterly Journal of Economics 122, 1351–1408

Blundell, R; Costa Dias, M; Meghir, C and Van Reenen, J (2004) 'Evaluating the Employment Impact of a Mandatory Job Search Assistance Programme', Journal of the European Economics Association, 2, 569-606.

Bond, S. and Cummins, J (2000) 'The Stock Market and Investment in the New Economy: Some Tangible Facts and Intangible Fictions', Brookings Papers on Economic Activity, 1, 61-124.

Borjas, G (2004) Labor Economics, McGraw-Hill (2nd edition).

Brown, C (1999) 'Minimum Wages, Employment, and the Distribution of Income', Chapter 32 in Orley Ashenfelter and David Card (eds.) Handbook of Labor Economics, North Holland Press.

Brown, C; Gilroy, C and Kohen, A (1982) 'The Effect of the Minimum Wage on Employment and Unemployment', Journal of Economic Literature, 20, 487-528.

Card, D (1992) 'Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage', Industrial and Labor Relations Review, 46,: 22-37.

Card, D and Krueger, A (1994) 'Minimum Wages and Employment: A Case Study of the Fast Food Industry in New Jersey and Pennsylvania', American Economic Review, 84, 772-793.

Card, D and Krueger, A (1995) Myth and Measurement: The New Economics of the Minimum Wage, Princeton University Press.

Culley, M; Woodland, S; O'Reilly, A and Dix, Gill (1999) Britain at Work, London: Routledge

Dickens, R; Machin, S and Manning, A (1999) `The Effects of Minimum Wages on Employment: Theory and Evidence From Britain', Journal of Labor Economics, 17, 1-22

DiNardo, J; Fortin, N and Lemieux, T (1996) 'Labor Market Institutions and the Distribution of Wages', Econometrica, 64, 1001-44.

DiNardo, J and Lemieux, T (1996) 'Changes in Wage Inequality in Canada and the United States: Do institutions explain the difference' Industrial and Labor Relations Review 50, 629-651.

Dickens, R and Manning, A (2004) 'Has the National Minimum Wage Reduced UK Wage Inequality?' Journal of the Royal Statistical Society, Series A, 167, 613-626.

Draca, M, Machin, S and Van Reenen, J (2008) 'Minimum Wages and Profitability', NBER Working Paper No. 13996

Fisher, F and McGowan, J (1983) 'On the misuse of the profits-sales ratio to infer monopoly profits' American Economic Review, 73, 82-97

Fougere, D; Erwan, G and le Bihan, H (2008) "Restaurant Prices and the Minimum Wage" CEPR Discussion Paper No. 6892

Freeman, R and Medoff, J (1984) What do Unions do?, New York: Basic Books.

Heckman, J; Ichimura, H and Todd, P (1997) 'Matching as an Econometric Evaluation Estimator', Review of Economic Studies, 65, 261-294.

Lemos, S (2008) 'A Survey of the Effects of Minimum Wages on Prices', Journal of Economic Surveys, 22, 182-212

Lemieux, T (2006) 'Increasing wage inequality: Composition effects, noisy data or rising demand for skill?' American Economic Review, 96(3), 461-98

Machin, S and Manning, A (1994) 'Minimum Wages, Wage Dispersion and Employment: Evidence from the U.K. Wages Councils', Industrial and Labor Relations Review, 47, 319-29.

Machin, S and Manning, A (2004) 'A Test of Competitive Labor Market Theory: The Wage Structure amongst Care Assistants in the South of England', Industrial and Labor Relations Review, 57, 371-85

Machin, S; Manning, A and Rahman, L (2003) 'Where the Minimum Wage Bites Hard: The Introduction of the UK National Minimum Wage to a Low Wage Sector', Journal of the European Economic Association, 1, 154-80.

Machin, S and Van Reenen, J (1993) 'Profit Margins and the Business Cycle: Evidence from UK Manufacturing Firms', Journal of Industrial Economics, 41, 29-50.

Machin, S and Wilson, J (2004) 'Minimum Wages in a Low Wage Labor Market: Care Homes in the UK', Economic Journal, 114, 102-109.

Metcalf, D (2002) 'The National Minimum Wage: Coverage, Impact and Future', Oxford Bulletin of Economics and Statistics, 64, 567-82.

Menezes-Filho, N (1997) 'Unions and Profitability Over the 1980s: Some Evidence on Union-Firm Bargaining in the United Kingdom', Economic Journal, 107, 651-70.

Metcalf, D (2003) 'Unions and productivity, financial performance and investment: international evidence' in Addison, John and Schnabel, Claus (eds), International Handbook of Trade Unions, London: Edward Elgar

Neumark, D and Nizolava, O (2007) 'Minimum Wages in the Longer Run', Journal of Human Resources, 42, 435-52.

Neumark, D and Wascher, W (1992) '"Employment Effects of Minimum and Sub-minimum Wage Laws", Industrial and Labor Relations Review, 46, 55-81.

Neumark, D and Wascher, W (2007) "Minimum Wages and Employment", Foundations and Trends in Microeconomics, 3, 1-182

Slade, M (2004) 'Models of Firm-Level Profitability', International Journal of Industrial Organization, 22, 289-308

Stewart, M (2002) 'Estimating the Impact of the Minimum Wage Using Geographical Wage Variation', Oxford Bulletin of Economics and Statistics, 64, 583-605.

Stewart, M (2004) 'The Impact of the Introduction of the UK Minimum Wage on the Employment Probabilities of Low Wage Workers', Journal of the European Economic Association, 2, 67-97.

Van Reenen, J (1996) 'The Creation and Capture of Economic Rents: Wages and Innovation in a Panel of UK Companies', Quarterly Journal of Economics, 111, 195-226.

Wadsworth, J (2009) 'Did the National Minimum Wage Affect UK Prices?', Centre for Economic Performance Discussion Paper No. 947.

3.8 Figures for Chapter 3

FIGURE 3.1: VALIDATION OF AVERAGE WAGE DATA (COMPARISON OF PROPORTION OF LOW WAGE WORKERS AND ESTABLISHMENT AVERAGE WAGES, WERS 1998)

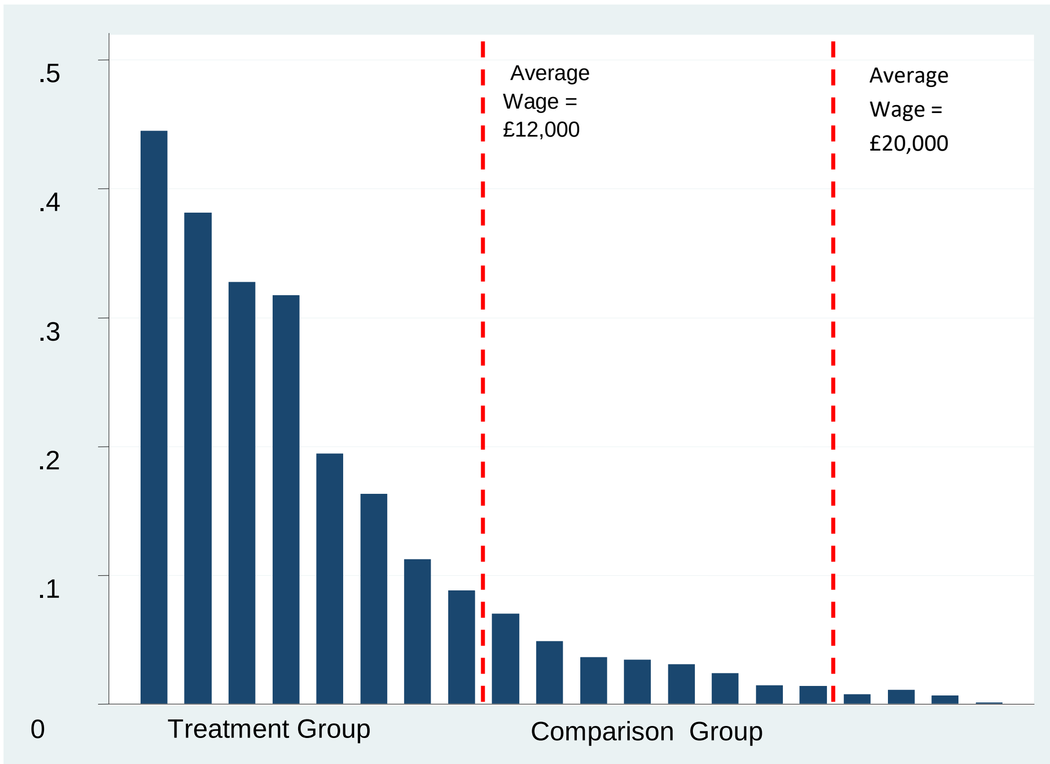

Notes:-
1). These figures are derived from the worker-establishment data (23,319 workers in 1,782 workplaces) from the 1998 Workplace Employee Relations Survey (WERS). The y-axis shows the proportion of workers paid below the minimum wage (£3.60 per hour) in the establishment. The x-axis shows the average annual wage at the workplace. This is divided into bins for of five percentiles from lowest (left) to highest (right) - a total of twenty bins up to an annualised wage of £24,000.
2). We mark the relevant thresholds for our analysis with vertical lines. The £12,000 line represents the main treatment group threshold used in our analysis of the FAME data. The £20,000 line is the cut-off for the upper bound of the comparison group used in the FAME analysis.

FIGURE 3.2: CHANGE IN LN(AVERAGE WAGE) BY PERCENTILE IN THE FINANCIAL YEAR BEFORE AND AFTER NMW INTRODUCTION

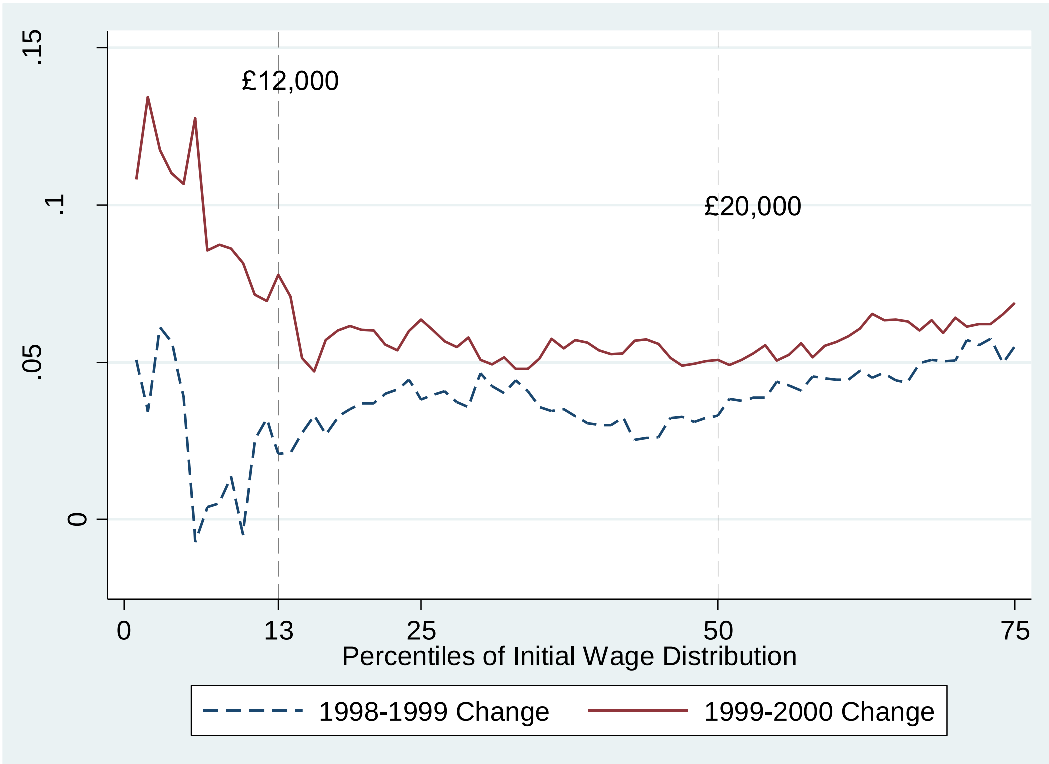

Notes:-

1). The data is taken from the FAME database of company accounts. The horizontal axis indicates the percentile in the firm wage distribution for a given firm in the initial period, the pre-policy financial year up to March 31st 1999. The vertical axis shows the proportionate change in average firm wages (between the pre-policy financial year and the post policy financial year) for each firm ranked by where it began in the wage distribution.

2). Pre-Policy is defined as the financial year April 1st 1998 to March 31st 1999; Policy On is defined as the financial year April 1st 1999 to March 31st 2000.

3) We show the threshold for the treatment groups by hatched vertical lines. In the baseline specifications firms with average wages below £12,000 (the 13th percentile) are in the treatment group and firms with average wages between £20,000 (the median) and £12,000 are in the control group.

FIGURE 3.3 VARYING TREATMENT EFFECT COEFFICIENTS IN FAME DIFFERENCE-IN-DIFFERENCE PROFITABILITY MODELS

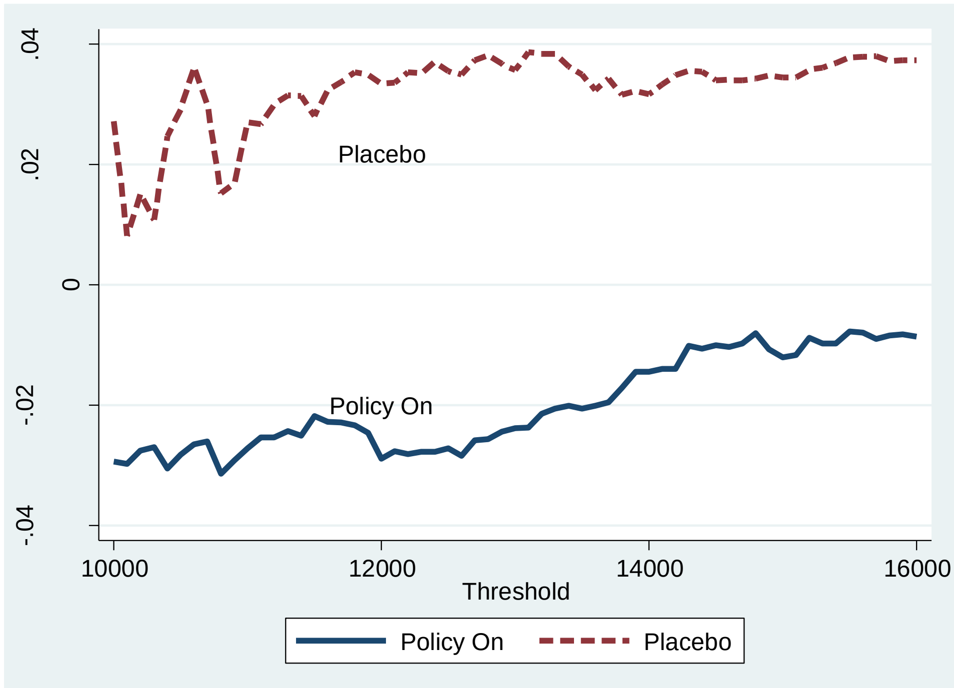

Notes:-

1). Data taken from is the FAME database of company accounts. The baseline models are as per Pre-NMW Low Wage Model in Tables 2 (Policy On period) and Table 3 (Pre-Policy period). .

2). The vertical axis shows the estimated treatment effects. The horizontal axis shows thresholds are shifted in units of £100 to define treatment group (T=1) as firms with pre-policy wages of under the threshold and comparison group with firms with average wages over the threshold and under £20,000. The baseline model is then re-defined and re-estimated using 50 successive treatment group wage thresholds between £10,000 and £15,000.

3). The Policy On sample period covers the six financial years from April 1st 1996 to March 31st 2002, NMW introduction on April 1st 1999. The Pre-Policy (pseudo-experiment) period covers the six financial years April 1st 1993 to March 31st 1999, with an 'imaginary' NMW introduction on April 1st 1996.

3.9 Tables for Chapter 3

TABLE 3.1: CHANGES IN FIRM AVERAGE WAGES AND PROFITABILITY BEFORE AND AFTER THE INTRODUCTION OF THE NATIONAL MINIMUM WAGE

	(1) Pre-NMW Introduction	(2) Post-NMW Introduction	(3) Difference
A. ln(Average Wage), lnW			
Pre-NMW Low Wage Firm, T=1	2.149	2.378	0.229
Pre-NMW Not Low Wage Firm, T=0	2.775	2.893	0.118
Difference-in-Difference			0.111*** (0.029)
B. Profit Margin, Π/S			
Pre-NMW Low Wage Firm, T=1	0.128	0.089	-0.039
Pre-NMW Not Low Wage Firm, T=0	0.070	0.058	-0.012
Difference-in-Difference			-0.027** (0.014)

Notes:-
1). Pre-NMW corresponds to the three financial years April 1st 1996-March 31st 1999 and Post-NMW to the three financial years April 1st 1999–March 31st 2002.
2). T = 1 indicates the treatment Group and T= 0 indicates the comparison group. Pre-NMW Low Wage Firm – the treatment group is defined as firms with an average wage equal to or below £12,000 per annum in the pre-policy financial year up to March 31st 1999; the comparison group is defined as firms with average wages between £12,000 and £20,000 in the pre-policy financial year up to March 31st 1999.
3). Standard errors in parentheses are clustered by firm and sample size is 4,112 (there are 951 firms).

TABLE 3.2: WAGES AND PROFITABILITY BEFORE AND AFTER INTRODUCTION OF THE NATIONAL MINIMUM WAGE (NMW), 1997-2002

	Period Before and After NMW Introduction, 1997-2002, (N = 4,112)	
	Change in ln(Average Wage), $\Delta \ln W$	*Change in Gross Profit Margin,* $\Delta(\Pi/S)$
A. Treatment = Low Wage Firm		
Pre-NMW Low Wage Firm	0.090*** (0.026)	-0.029** (0.012)
Test of no behavioral response	P-value = 0.663	
B. Treatment = - Pre-Policy ln(W)		
- Pre-NMW ln(W)	0.188*** (0.033)	-0.032** (0.015)
Test of no behavioral response	P-value = 0.144	

Notes:-
1). Coefficients estimated by Ordinary Least Squares and standard errors in parentheses below are clustered by firm (there are 951 firms).
2). The Pre-NMW period covers the three pre-policy financial years April 1st 1996-March 31st 1999 and the Post-NMW period the three financial years April 1st 1999–March 31st 2002. Low Wage Firm Pre-NMW - treatment group is defined as firms with an average wage equal to or below £12,000 per annum in the pre-policy financial year up to March 31st 1999; the comparison group is defined as firms with average wages between £12,000 and £20,000. Pre-NMW ln(W) - indicates that we use a continuous measure of the wage (in the pre-policy year up to March 31st 1999) is used for treatment intensity.
3). Controls include two digit industry dummies; 18 regional dummies; the proportion of workers who are graduates (by region and two-digit industry); and union membership, part-time work and female employment rates (by three-digit industry classification).
4). “Test of no behavioral response” implements equation (3) in the text.

TABLE 3.3: WAGES AND PROFITABILITY BEFORE AND AFTER INTRODUCTION OF A PLACEBO NATIONAL MINIMUM WAGE (NMW), 1993-1999

	Period Before and After 'Imaginary NMW' Introduction, 1993-99, (N = 4,550)	
	Change in ln(Average Wage), $\Delta \ln W$	*Change in Gross Profit Margin,* $\Delta(\Pi/S)$
A. Treatment = Low Wage Firm		
Pre-'Imaginary NMW' Low Wage Firm	0.033 (0.028)	0.015 (0.011)
B. Treatment = - Pre-Policy ln(W)		
- Pre-'Imaginary NMW' ln(W)	0.079 (0.106)	0.012 (0.029)

Notes:-
1). Coefficients estimated by Ordinary Least Squares and standard errors in parentheses below are clustered by firm (there are 1,047 firms).
2). The Pre-'Imaginary NMW' period covers the three financial years April 1st 1993 – March 31st 1996 and the Post-'Imaginary NMW' period covers the three financial years April 1st 1996 and March 31st 1999. Low Wage Firm Pre-'Imaginary NMW' - treatment group is defined as firms with an average wage equal to or below £12,000 per annum in the pre-policy financial year up to March 31st 1996; the comparison group is defined as firms with average wages between £12,000 and £20,000. Pre-'Imaginary NMW' ln(W) - indicates that we use a continuous measure of the wage (in the Pre-'Imaginary NMW' year up to March 31st 1996) is used for treatment intensity.
3). Controls include two digit industry dummies; 18 regional dummies, the proportion of workers who are graduates (by region and two-digit industry); and union membership, part-time work and female employment rates (by three-digit industry classification).

TABLE 3.4: NATIONAL MINIMUM WAGE INTRODUCTION AND WAGES AND PROFITABILITY IN CARE HOMES, 1998-1999.

	Period Before and After NMW Introduction, 1998-99, (N = 908)	
A. Wages	*Change in ln(Average Wage),* $\Delta \ln W$	
Pre-NMW Wage Gap	0.861*** (0.045)	0.886*** (0.052)
Controls	No	Yes
B. Profitability	$\Delta(\Pi/S)$, *Change in Profit Margin*	
Pre-NMW Wage Gap	-0.433*** (0.173)	-0.492*** (0.202)
Controls	No	Yes

Notes:-
1). Coefficients estimated by Ordinary Least Squares. Robust standard errors in parentheses under coefficients.
2). Sample covers 454 nursing homes in 1998 and 1999.
3). Initial pre-minimum wage period (t-1) controls include workforce characteristics (proportion female, mean worker age, proportion with nursing qualifications), the proportion of residents paid for by the government ("DSS"), region dummies and month dummies.

TABLE 3.5: SPLITTING INTO HIGH AND LOW MARKET POWER INDUSTRIES

Outcome	High Market Power Industries	Low Market Power industries
A. Wages		
Treatment = Low Wage Firm	0.109***	0.081**
N = 1,943 (High); N =2,169 (Low)	(0.035)	(0.038)
B. Profits		
Treatment = Low Wage Firm	-0.037**	-0.014
N = 1,943 (High); N =2,169 (Low)	(0.018)	(0.014)
Test of no behavioral response	P-value = 0.646	P-value = 0.531
C. Employment		
Treatment = Low Wage Firm	0.104	-0.012
N = 1,943 (High); N =2,169 (Low)	(0.142)	(0.121)
D. Labor Productivity		
Treatment = Low Wage Firm	0.075	0.113
N = 1,943 (High); N =2,169 (Low)	(0.066)	(0.090)
E. Exit		
Treatment = Low Wage Firm	-0.023	-0.002
N=1,150 (High); N= 1,206 (Low)	(0.023)	(0.027)

Notes:-
1). This table shows the results from a series of separate regressions for the Low Wage Firm models. The dependent variable is indicated in the first column, Column (1) is on the sub-sample of firms in high market power industries and column (2) is the sub-sample of firms in the low market power industries
2). High market power industries are defined as those with higher than the median value of the industry-level Lerner Index in the firm's three-digit industry. Low market power industries are defined as those with below the median value of the industry-level Lerner Index in the firm's three-digit industry.
3). Coefficients estimated by Ordinary Least Squares and standard errors in parentheses below are clustered by firm.
4). Employment is total number of workers in the firm. Labor productivity is ln(sales/employment). "Exit" is defined for two cohorts in 1996 (pre-NMW) and 1999 post NMW and indicates whether the firm ceased to exist in the subsequent 3 years (see text)
5). Controls include two digit industry dummies; 18 regional dummies, the proportion of workers who are graduates (by region and two-digit industry); and union membership, part-time work and female employment rates (by three-digit industry classification).

TABLE 3.6: FIRM ENTRY AND EXIT (BY THREE-DIGIT INDUSTRY).

	Period Before and After NMW Introduction, 1996-2001, (N = 1,020)	Period Before and After 'Imaginary NMW' Introduction, 1994-98, (N = 850)	Difference
A. Change in Industry Entry Rates			
Pre-NMW Low Pay Proportion	0.021 (0.015)	0.057* (0032)	-0.036 (0.038)
B. Change in Industry Exit Rates			
Pre-NMW Low Pay Proportion	-0.013 (0.016)	-0.028 (0.018)	0.015 (0.024)
C. Change in Industry Net Entry Rates			
Pre-NMW Low Pay Proportion	0.034 (0.025)	0.085** (0.027)	-0.051 (0.037)

Notes:-
1). Data taken from Value-Added Tax (VAT) Registrations and Deregistration Data, Department of Trade and Industry (DTI). Entry rate is the proportion of firms who are newly registered in a year in a three-digit industry. Exit rate is the proportion of firms who are deregistered in the year. Net entry is entry rate – exit rate. Standard errors (in parentheses) are clustered by three-digit industry.
2). Pre-NMW low pay proportion is the proportion of workers with an hourly wage less than £3.60 in the three-digit industry in real terms over the pre-policy period (the minimum wage threshold of £3.60 is deflated by the retail price index for the years 1994-1998).
3). All specifications include controls for two digit industry dummies, time dummies, and the proportion of employees in the three-digit industry that are female, part-time and the proportion of employees in the three-digit industry that are female, part-time and unionized.

3.10 Appendix for Chapter 3

3.A1. Introduction

In order to obtain a long-run effect of the minimum wage on profitability we need to have some degree of imperfect competition in the product market. We therefore consider several industrial organization models. Aaronson and French (2007) consider in detail the effects on the minimum wage on prices and costs in a competitive and monopsonistic labor market model. However, in these models firms do not have positive price-cost margins so profits remain zero by assumption, regardless of the minimum wage.

We separate our analysis into the short-run and long-run, where we define the short run as the period where all variables are able to change (*including* capital, labor, prices, etc) but the number of firms is held fixed. In the long-run entry and exit can occur and the number of firms can change. Our analysis of exit and entry is directly applicable to the long-run results.

3.A2. Imperfect Competition in the Product market

Short-run effects with symmetric and asymmetric costs

Consider a two-stage game where firms pay a sunk entry cost (K) and, conditional on entering engage in competition with other firms (total number of firms in market is denoted N). The instruments of competition can be price or quantity.

We begin with the workhorse industrial organization model of an asymmetric Cournot model[69] where firms have heterogeneous marginal costs. Below we discuss alternative imperfect competition models that lead to similar qualitative results.

The non-cooperative Nash equilibrium in quantities gives a well-know expression for the price-cost margin:

$$\frac{p - c_i(q_i)}{p} = \frac{MS_i}{\eta} \tag{A1}$$

---

[69] Cournot competition can be considered the reduced form of a two-stage game where firms set capacities in the first stage and then compete in prices in the first stage (Kreps and Scheinkman, 1983).

Where η is the (absolute value of the) price elasticity of product demand, p is output price, c_i is marginal cost of firm i, q_i is firm output and MS_i is the market share ($\mathrm{MS}_\mathrm{i} = \frac{S_i}{\sum_i S_i}$) with S_i denoting firm sales. Note that equation (A1) nests the special cases of monopoly (N = 1). If we assume constant returns then marginal costs do not depend on output ($c_i = c_i'(q_i)$) so the price-cost margin can be characterised by the ratio of profit (Π) to sales (S):

$$\left(\frac{\Pi}{S}\right)_i = \frac{MS_i}{\eta} \tag{A2}$$

Firm i's market share will depend on its marginal costs relative to the marginal costs of other firm's in the industry. If firm i's marginal costs rise relative to those of other firms it will lose market share (see Tirole, 1989, Chapter 5 for example).

Consider the effect of an increase in the minimum wage. If we assume that demand is iso-elastic (we will relax this below) then the impact of the minimum wage on the firm's price-cost margins will be reflected in its market share. If a firm employs a greater proportion of minimum wage workers, it will face a larger increase in marginal costs and therefore a larger fall of its price-cost margins.

This is our key comparative static result: the introduction of a minimum wage will reduce the profitability of firms who are more "at risk" because they employ a higher share of minimum wage workers.

If we also relax the assumption the demand elasticity is constant, there will also likely be a fall in profitability. To see this clearly assume that firms are symmetric so that they all face identical marginal costs. In this case, the equilibrium condition of (A1) simplifies to

$$\frac{\Pi}{S} = \frac{1}{N}\frac{1}{\eta} \tag{A3}$$

It is clear from equation (A3) that the impact of the minimum wage will on profitability ($\frac{\Pi}{S}$) will depend on its impact on the demand elasticity (η). In particular if demand becomes more elastic, profitability will fall. For most commonly used demand curves, a minimum wage will make the demand curve more elastic because price has risen. For example, consider the case of linear industry demand (Q) for where $Q = A - bp$, $b > 0$, $A > 0$. In this case, $\eta = b\frac{p}{Q}$.

Following the introduction of a minimum wage prices will be higher and quantity sold lower unless demand is perfectly elastic. The elasticity of demand is therefore higher and profitability will fall. This will reinforce the effects on market share discussed in the more general model with asymmetric firms.[70]

Under differentiated products Bertrand equation (A3) should be interpreted as a firm-specific elasticity. This result differs from Aaronson and French (2006) who consider a model of monopolistic competition. This generates an equilibrium condition like (A3). The minimum wage has no effect on price cost margins in their model because they assume that the elasticity of demand is constant. This guarantees no effect of the minimum wage on price-cost margins as all costs are passed through completely to the consumer. Additionally, the "large number of firms" assumption underlying monopolistic competition rules out strategic interactions that generate the market share effects in equation (A3)

Long-run effects

After the minimum wage is imposed, absolute profits in the industry will be lower. This will mean that there is less of an incentive to enter the industry. Consequently, we might expect to see fewer firms in the industry (from exit and/or less entry) in the long run. The short run fall in profits for the incumbent firms in the industry will therefore be greater than the long-run change as N will fall (e.g. see equation (A3)).

An important caveat to this is that the number of firms in the industry may not fall due to an "integer" effect. Since there will always be an integer number of firms in the industry all firms will usually earn some economic profit. Firms will enter and pay the sunk cost up until the point that a marginal firm entering the industry would not make a profit net of the sunk cost. For

[70] We cannot rule out the possibility that the aggregate demand curve may become more elastic as wages rise even if the labor market is perfectly competitive. Micro-economic theory places few restrictions on industry demand curve aggregated from consumer preferences (e.g. see Varian, 1984, chapter 3.16). Thus, it is still ultimately an empirical issue whether profitability rises or falls after the minimum wage.

example, consider a symmetric duopoly in long-run equilibrium. If a third firm entered the industry a firm's profits (net of the sunk cost, *K*) would be negative i.e. $\Pi^{*(3)} - K \leq 0 \leq \Pi^{*(2)} - K$, where $\Pi^{*(3)}$ is equilibrium profits with three firms and $\Pi^{*(2)}$ equilibrium profits with two firms.

Now, except in the special case when profits in the market exactly covers the sunk cost ($\Pi^{*(3)} - K < 0$ and $\Pi^{*(2)} - K = 0$) the minimum wage could reduce $\Pi^{*(2)}$, but not by so much that $\Pi^{*(2)} - K$ <0 and one firm was forced to exit the industry. Consequently, for small increases in the minimum wage firms could have lower profits without a change in the equilibrium number of firms.

This caveat aside, in a dynamic setting we would expect that a minimum wage would increase exit and reduce the entry rate.

3A3. Perfect Competition in the Product Market

Now consider the case of perfectly competitive product markets. Comparative statics of prices and factor demands following a minimum wage increase have been comprehensively analyzed by Aaronson and French (2007). Here, we will briefly contrast the usual case of perfect competition in the labor market with some alternative models. It worth emphasizing two preliminary points. First, as we discussed above that these are in the short-run as in the long run firms earn zero profits by assumption. Second, the short-run effect of the introduction of a minimum wage will be *larger* in the competitive model than in the monopsony model.

Perfect competition in the labor market

If labor markets are perfectly competitive, the short run effects of the minimum wage on profits are composed of two components (see Ashenfelter and Smith, 1979, and the main text). First, there is fall in profits due to the increased wage for the current number of workers paid below the minimum wage. This fall in profits is offset by a second effect to the degree that firms can substitute minimum wage workers for other factors of production (including non minimum wage workers). In the limiting case of perfect substitutability of minimum wage workers there will be *no effect* on profits.

Of course these are only short-run effects as there can be no economic profits under perfect competition and in equilibrium industry prices will rise and quantity will fall (so there will either be fewer firms or the average firm size will shrink).

Imperfect Competition in the labor market

There have been a variety of models proposed in recent years where firms have some power to set wages because of efficiency wages, monopsony, search or other reasons. In these models, over a certain range of values a binding minimum wage can increase employment.

Considering profits, we would expect the negative short-run effects of a minimum wage on profitability to be muted in such models. This is because, unlike the competitive model the first order effect on profits is zero as an increase in the wage has a beneficial effect on profits through making it easier to recruit, retain and/or motivate workers. There will be a second effect because the firm is being shifted away from its optimal level of the wage so overall we would still expect a decline in profits. However, this is likely to be much less severe than in the competitive model.

To see this consider a simple representation of the monopsony model. We model the firm's wage setting power in a reduced form way (following Card and Krueger, 1995) and assume that the production function, $F(W,L)$, is increasing in the wage as well as labor, L. The firm chooses wages and labor to maximize profits

$$\Pi = \max_{W,L}\ pF(W,L) - WL$$

which lead to the standard first order condition:

$$p\frac{\partial F(W,L)}{\partial L} = W^*$$

where an asterix denotes the optimized value. We also have an additional non-standard condition from optimizing wages of:

$$p\frac{\partial F(W,L)}{\partial W} = L^* \qquad \text{(A4)}$$

If we consider the effect of a small increase in wages on profits in the neighbourhood of the optimized level of wages and employment (W^*, L^*) this is given by:

$$\frac{d\Pi}{dW} = p\frac{\partial F(W^*, L^*)}{\partial W} - L^*$$

Note that this is equal to zero by the first order condition with respect to wages, equation (A4).

Long-run effects

In this setting, there are no long-run effects on profits. Considering exit, unlike the model with imperfect competition firm size is not tied down. In the competitive model, prices will be higher and output lower. In our constant returns set-up a zero profit equilibrium can be restored either by all firms becoming smaller or by some firms exiting.

3.A4. Summary

In models of imperfect product market competition, we would generally expect to observe negative effects on the profitability of firms where the minimum wage bites, even after firms have adjusted all factors of production. In such models, some of the increase in costs is borne by shareholders rather than just consumers and unemployed low-wage workers as in the standard competitive model. It is worth emphasizing that employment will still fall in these models. So oligopoly could explain only why employment responses could be more muted than one would expect from a competitive model. Of course, employment changes can be positive if firms with market power in the product market also have market power in the labor market.

The final section (A4) showed a very simple model that assumes no change in sales or jobs following a minimum wage hike. This model does surprisingly well in rationalizing the results.

Appendix 3.B:Data

FAME Data

The FAME (Financial Analysis Made Easy) dataset contains information on firm company accounts of publicly listed and private firms in the UK economy. It is supplied under licence as part of the AMADEUS database from BVD (Bureau Van Dijk). Our sample begins with data on all firms in the six financial years from April 1st 1996 to March 31st 2002 including those who had entered and exited. We select firms who report on the 31st March. We drop firms with missing data on our key variables (profits, wages, sales, employment, industry, and region). We use information on consolidated accounts at the lowest level that exists (i.e. we use subsidiary level information if this exists). We drop information for all observations where the profit-sales ratio is greater than 1 in absolute value.

In the main results, we condition on the cohort of firms who were alive on March 31st 1999 when the minimum wage introduced and had an average wage between £4000-£20,000. We also present results where we examine the impact of including firms who entered after this date (and exited before this date) including a dummy variables for entrant and exiting firms (and interactions of these dummies with the NMW policy period).

Profits/Sales: Gross profits (prior to deductions for tax, interest and dividends) over turnover (sales).

Average Wages: Total remuneration divided by total number of employees

Capital / Sales: Tangible assets over turnover (sales).

Sales / Employment: Total turnover (sales) over the number of employees.

Labor Force Survey

The Labor Force Survey (LFS) is a large-scale household interview-based survey of individuals in the UK that has been carried out on varying bases since 1975.[71] Around 60,000

[71] Between 1975 and 1983, the survey was conducted every two years. From 1984 until 1991, it was conducted annually. Since 1992, the Labor Force Survey has been conducted every three months in a five-quarter rolling panel format.

households have been interviewed per survey since 1984. Annual proportions calculated relative to firm reporting year rather than calendar year (i.e. April 1998 – March 1999).

Union membership: Defined at the three-digit UKSIC industry level, annual values 1993-2002.

Part-Time Work: Proportion of employed workforce classified as part-time, annual values 1993-2002. Defined at the three-digit industry level.

Female Workforce: Female workers as a proportion of total employed workforce, annual values 1993-2002. Defined at the three-digit industry level.

Graduate Qualifications: Proportion of graduate qualified workers per region and two-digit industry cell.

Region: Government Office Region of Workplace ("gorwk"). These include Tyne and Wear, Rest of the North East, Greater Manchester, Merseyside, Rest of the North West, South Yorkshire, West Yorkshire, Rest of Yorkshire and Humberside, West Midlands and Met Country, Rest of West Midlands, Eastern, Inner London, Outer London, South East, South West, Wales, Rest of Scotland, Northern Ireland.

Care Homes Data

The UK care homes data was collected in surveys conducted in 1992 (prior to the general election in that year) and 1993 for homes on the South Coast of England; in 1998 (before the introduction of the NMW) and in 1999 (after the introduction of the NMW in April) for all homes across the country. Finally, there was some more data collected in 2000 and 2001 for South Coast homes only. The data is in the form of an unbalanced panel so that the same homes are followed over time. The sector was chosen because it is characterized by a large concentration of non-unionized, low wage employees working in small firms with an average employment level of fifteen to twenty. There was also product market regulation in this sector insofar that an important fraction of home residents had their care paid for by the government through the Department of Social Security (DSS).[72] The Department of Social Security paid a capped price for beds, which were not increased when the minimum was introduced. As a result,

[72] The average percentage of such residents was 52.7% before the minimum wage introduction and 57.6% after. We always condition on this variable in the regressions.

many homes had a limited scope to increase prices in response to the minimum thereby leaving more room for employment or profitability effects to manifest themselves. A more comprehensive account of features of the data is given in Machin, Manning and Rahman (2003).

Business Registration and De-registration Database

The UK Department of Trade and Industry (DTI) publish data on births and deaths of companies at the three-digit level on a consistent basis from 1994 (see http://stats.berr.gov.uk/ed/vat/). These are based on Value Added Tax (VAT) Registration numbers that every incorporated firm in Britain is legally obliged to have. (This is the same as the aggregated FAME date).

We used this data to calculate for each three-digit sector the proportion of firms who entered in a year (entry rate). Entry rates calculated as the number of new VAT (Value-Added Tax) registrations as a proportion of the beginning of year stock. Exit rate calculated as the number of VAT deregistrations over the beginning-of-year stock. Net entry calculated as entry rate minus exit rate. We also calculated the net entry rate as the difference between the entry and exit rates.

We then matched information form the LFS at the same level of aggregation to calculate the proportion of workers in each industry paid below the minimum wage in the pre-policy period.

TABLE B1:
CHARACTERISTICS OF TREATMENT AND COMPARISON GROUPS

	Treatment Group T=1	Comparison Group T=0	All
Average Wage (£000s)	10.53	17.38	15.76
Profit/Sales	0.108	0.064	0.074
Capital/Sales	0.297	0.237	0.248
Wagebill/Sales	0.289	0.261	0.268
Employment (mean)	2,704	1,004	1,407
Employment (median)	273	170	187
Productivity (=Sales/Employee) (£000s)	71.4	110.2	101.0
Exit Rate	0.050	0.053	0.051
Proportion part-time employees	0.295	0.158	0.190
Proportion female employees	0.535	0.378	0.415
Proportion union members	0.186	0.213	0.207
Proportion Firms in:			
Manufacturing	0.165	0.372	0.323
Wholesale	0.081	0.172	0.150
Retail	0.098	0.038	0.052
Hospitality	0.163	0.015	0.050
Business Services	0.133	0.083	0.095
Number of Observations	974	3,138	4,112

NOTES:- T= 0: Comparison group; T = 1: Treatment Group; Part-time and female employees based on Labor Force Survey (LFS) and calculated as proportion of total workers per two-digit industry by regional cell. *Low Wage Firm* - treatment group is defined as firms with an average wage equal to or below £12,000 per annum in the pseudo pre-policy financial year up to March 31st 1996; the comparison group is defined as firms with average wages between £12,000 and £20,000. Sample for exit represents 1999 cohort of firms, with total N = 1,066 (N=319 for treatment group and N=747 for comparison group).

TABLE B2: FIRM ENTRY AND EXIT RATES BY THREE-DIGIT INDUSTRY, 1996-2001 (DTI VAT REGISTRATIONS AND DEREGISTRATIONS).

	(1) All Industries	(2) Low Wage industries (below median Lowpay)	(3) High Wage Industries (above median Lowpay)
Entry Rate	0.089	0.087	0.091
Exit Rate	0.082	0.083	0.081
Net Entry	0.007	0.003	0.011
Lowpay	0.126	0.051	0.201
Union	0.287	0.350	0.189
Female	0.343	0.274	0.411
Part-time	0.143	0.076	0.209
No. of Industries	170	85	85
No. of Observations	1,020	510	510

NOTES: Entry rates calculated as the number of new VAT (Value-Added Tax) registrations as a proportion of the beginning of year stock. Exit rate calculated as the number of VAT deregistrations over the beginning-of-year stock. Net entry calculated as entry rate minus the exit rate. The variables lowpay, union, female, part-time are all sourced from the UK Labor Force survey (LFS). The “Lowpay” variable is defined as the proportion of workers with hourly wage below £3.60 in the pre-minimum wage period (1994-1998). “Below Median Lowpay” indicates all those industries where the proportion of lowpay workers ranges from 0 to 0.092. “Above Median Lowpay” indicates all of those industries where the proportion of lowpay industries ranges from 0.095 to 0.557.

CHAPTER 4: TRADE-INDUCED TECHNICAL CHANGE? THE IMPACT OF CHINESE IMPORTS ON INNOVATION, IT AND PRODUCTIVITY

Abstract

We examine the impact of Chinese import competition on broad measures of technical change - patenting, IT, R&D, TFP and management practices – using new panel data across twelve European countries from 1996-2007. In particular, we establish that the *absolute* volume of innovation increases within the firms most affected by Chinese imports. We correct for endogeneity using the removal of product-specific quotas following China's entry into the World Trade Organization. Chinese import competition led to increased technical change *within firms* and reallocated employment *between firms* towards more technologically advanced firms. These within and between effects were about equal in magnitude, and account for about 15% of European technology upgrading over 2000-2007 (and even more when allowing for offshoring to China). Rising Chinese import competition also led to falls in employment, profits, prices and the share of unskilled workers. By contrast, import competition from developed countries had no effect on innovation. We develop a simple "trapped factor" model that is consistent with these empirical findings.

JEL Code. O33, F14, L25, L60,

Keywords: China, technical change, trade, firm survival, employment

4.1 Introduction

A vigorous political debate is in progress over the impact of globalization on the economies of the developed world. China looms large in these discussions, as her exports have grown by over 15% per year over the last two decades. One major benefit of Chinese trade had been lower prices for manufactured goods. We argue in this paper that increased Chinese trade has also induced faster technical change from both innovation and the adoption of new technologies, contributing to productivity growth. In particular, we find that the *absolute* volume of innovation (not just patents per worker or productivity) increases *within* the firms more affected by exogenous reductions in barriers to Chinese imports.

Several detailed case studies such as Bartel, Ichinowski and Shaw (2007) on American valve-makers, Freeman and Kleiner (2005) on footwear or Bugamelli, Schivardi and Zizza (2008) on Italian manufacturers show firms innovating in response to import competition from low wage countries. A contribution of our paper is to confirm the importance of low wage country trade for technical change using a large sample of over half a million firms and exploiting China's entry into the World Trade Organization (WTO) to identify the causal effect of trade.

The rise of China and other emerging economies such as India, Mexico and Brazil has also coincided with an increase in wage inequality and basic trade theory predicts such South-North integration could cause this. Despite this, the consensus among most economists was that trade was less important than technology in explaining these inequality trends.[73] There are three problems with this consensus, however. First, most of this work used data only up to the mid 1990s, which largely predates the rise of China (see Figure 4.1).[74] Note that we end our sample in 2007 prior to the Great Recession to avoid conflating the effect of China with that of the financial crisis and subsequent huge fall in trade. Second, Feenstra and Hansen (1999) points to an impact of trade through offshoring rather than final goods. Third, an emerging line of theory has pointed to mechanisms whereby trade can affect the incentives to develop and adopt new technologies. Thus, the finding that measures of technical change are highly correlated with skill upgrading does not mean trade has no role. What may be happening is that trade is stimulating technical progress, which in turn is increasing the demand for skilled labor.

Our paper addresses these three problems. First, we use data from the last decade to examine the recent role of trade in affecting technical change in developed countries. Second, we will examine offshoring to China. Third, we analyze the impact of imports on patents, information technology (IT), research and development (R&D), total factor productivity (TFP) and management practices. We distinguish between the impact of import competition on

[73] See, for example, Acemoglu (2002), Autor, Katz and Kruger (1998), Machin and Van Reenen (1998) and DiNardo, Fortin, and Lemieux (1996).

[74] In the 1980s China only accounted for about 1% of total imports to the US and EU and by 1991 the figure was still only 2%. However, by 2007 China accounted for almost 11% of all imports and Krugman (2008) emphasises this in his re-evaluation of the older literature. Note that Figure 1 may overestimate China's importance as import growth does not necessarily reflect value added growth. For example, although IPods are produced in China, the intellectual property is owned by Apple. However, our identification relies on *differences* in Chinese imports over time and industries, and our results are stronger when we use quota abolition as an instrumental variable, so using import value (rather than value added) does not appear to be driving our results.

technology through a within firm effect and a between firm (reallocation) effect and find that *both* matter.

A major empirical challenge in determining the causal effect of trade on technical change is the presence of unobservable technology shocks. To tackle this endogeneity issue we implement three alternative identification strategies. Our main approach is to use China's entry into the World Trade Organization (WTO) in 2001 and the subsequent elimination of most quotas in the ensuing years under the Agreement on Clothing and Textiles (formerly the Multi Fiber Agreement). These sectors are relatively low tech, but were still responsible for over 22,000 European patents in our sample period. Second, we exploit the fact that the exogenous liberalization policies in China had differential effects on imports into Europe across industries. In particular, Chinese import growth in Europe was much stronger in the sectors where China had some comparative advantage. Third, we control for differential industry-specific time trends. All three identification strategies support our main finding that Chinese trade stimulates faster technical change.

We present two core results. First, on the intensive margin, Chinese import competition increases innovation, TFP and management quality *within* surviving firms. Firms facing higher levels of Chinese import competition create more patents, spend more on R&D, raise their IT intensity, adopt more modern management practices, and increase their overall level of TFP. Second, Chinese import competition reduces employment and survival probabilities in low-tech firms - e.g. firms with lower levels of patents or TFP shrink and exit much more rapidly than high-tech firms in response to Chinese competition. Thus, our paper jointly examines the effects of trade on survival/selection and innovation. The combined impact of these within firm and between effects is to cause technological upgrading in those industries most affected by Chinese imports. An additional set of results shows that Chinese imports significantly reduce prices, profitability and the demand for unskilled workers as basic theory would suggest.

We focus on China both because it is the largest developing country exporter, and because China's accession to the WTO enables us to plausibly identify the causal effects of falling trade barriers. However, we also show results for imports from all other developing countries, and find a similar impact on technical change. In contrast, imports from developed countries appear to have no impact on technology.

We also offer some back of the envelope quantification of Chinese import effects on technical change. Over 2000-2007 China appeared to account for almost 15% of the increase in patenting, IT and productivity. Furthermore, this effect has grown in recent years and is up to twice as large when incorporating offshoring. These results suggest that trade with emerging nations such as China may now an important factor for technical change and growth in richer countries.

To motivate the empirical framework we discuss a model, further developed in Bloom, Romer, Terry and Van Reenen (2012), that explains how trade from China drives innovation in exposed firms. The intuition relies on "trapped-factors" – that is factors of production which are costly to move between firms because of adjustment costs and sunk investment (e.g. firm-specific skills and capital). Chinese imports reduce the relative profitability of making low-tech products but since firms cannot easily dispose of their "trapped" labor and capital, the shadow cost of innovating has fallen. Hence, by reducing the profitability of current low-tech products and freeing up inputs to innovate, Chinese trade reduces the opportunity cost of innovation. In addition to Chinese import competition stimulating innovation, we find support for two other predictions of the model. First, import competition from low wage countries like China has a greater effect on innovation than imports from high wage countries. This occurs because Chinese imports have a disproportionate effect on the profitability of low-tech products, providing greater incentives to innovate new goods. Second, firms with more trapped factors (as measured by industry-specific human capital, for example) will respond more strongly to import threats.

Our paper relates to several literatures. First, for labor economics we find a role for trade with low wage countries in increasing skill demand (at least since the mid-1990s) through inducing technical change.[75] Second, although many papers have found that trade liberalization increases aggregate industry productivity[76], the precise mechanism is unclear. This evidence tends to be indirect since explicit measures of technical change are generally unavailable at the micro-level.[77] The literature focuses on the reallocation effects (e.g. Melitz, 2003) even though within plant productivity growth is typically as large as the between-plant reallocation effect. Our paper uses new patenting, IT, R&D, management and productivity data

---

[75] Technological forces also have an effect on trade. For example, better communication technologies facilitate offshoring by aiding international coordination. This is another motivation for addressing the endogeneity issue. Additionally, there is the direct impact on local employment and welfare (e.g. Autor, Dorn and Hansen, 2012).

[76] See, for example, Pavcnik (2002), Trefler (2004), Eslava, Haltiwanger and Kugler (2009), and Dunne, Klimek and Schmitz (2008).

[77] For low-wage countries, Bustos (2011) finds positive effects on innovation from lower export barriers for Argentinean firms and Teshima (2008) finds positive effects on process R&D from lower output tariffs for Mexican firms. The only study of Southern trade on Northern innovation is LeLarge and Nefussi (2008), who find that the R&D of French firms reacts positively to low wage country imports, although they have no external instrument.

to establish that trade drives out low-tech firms (reallocation) and increases the incentives of incumbents to speed up technical change.

Third, there is a large theoretical literature on trade and technology.[78] Our paper supports theories arguing for an important role of trade on technical change. In particular, our finding that (i) the positive trade effect is on *innovation* (rather than just compositional effects on productivity via offshoring or product switching) and (ii) is much stronger from lowering import barriers against low-wage countries rather than high-wage countries is different from the mechanisms emphasized in other theories (e.g. market size or learning).

Finally, there is a large empirical literature examining the impact of competition on innovation, but a major challenge is finding quasi-experiments to identify the causal impact of competition on innovation (e.g. Aghion et al, 2005). Our paper extends this work by using China's trade growth, and particularly its entry into the WTO, as an exogenous shift in competition. The structure of the paper is as follows: Section 4.2 sketches some theoretical models, Section 4.3 describes the data and Section 4.4 details the empirical modeling strategy. Section 4.5 describes our results and Section 4.6 discusses their magnitudes. Some extensions and robustness tests are contained in Section 4.7 and Section 4.8 concludes.

4.2 Theory

There are a large number of theories of how reducing import barriers against low wage countries (like China) could affect technical change in high wage countries (like Europe or the US). We first outline a simple "trapped factor" model that predicts a positive effect of such

[78] Theoretical analysis of trade and innovation is voluminous from the classic work by Grossman and Helpman (1991, 1992) and recent important contributions by Yeaple (2005) and Atkeson and Burstein (2010).

liberalization on *innovation*, and two ancillary predictions. We contrast this with other perspectives on innovation (where trade expands the menu of products in the world economy) and *composition* where trade alters the distribution of products without changing the number of quality of products.

4.21 The "Trapped Factor" model of Trade-induced innovation

In Bloom, Romer, Terry and Van Reenen (2012) we develop a stylized model of trade-induced innovation (see Appendix A for a more detailed summary). The basic assumption is that firms can allocate a factor of production either to produce old goods or innovate and produce new goods. China can produce old goods, but cannot (as easily) innovate and produce new goods. At the beginning of the period there are factors of production employed in "Northern" firms making old goods (protected by trade barriers). These factors are "trapped" in the sense that there is some human or fixed capital that is specific to the old good that is lost for a period if the firm chooses to reallocate the factor from producing the old good to innovating a new good. The magnitude of the firm-specific capital determines the opportunity cost of innovation and if it is sufficiently high the firm optimally chooses not to innovate.

When import barriers are lowered, Chinese exports increase and the profitability of making old goods falls. Therefore, the opportunity cost of using the trapped factors for innovating (rather than producing the old good) falls, making innovation more attractive[79]. Not

[79] In the model we make the simplifying assumption that the firms who innovate also produce the good while it is on patent. When it comes off patent the good is produced in the home country if protected by trade barriers or in the South if not protected. If we extend the model to allow for offshoring, then the innovation could still occur in the rich country but production of the new good could occur in the low wage country (e.g. the R&D for the I-Pod is in the US, but it is produced in China). In this case, reducing trade barriers with China will make the costs of *producing* the new good cheaper and this could be an additional reason why incentives to innovate on new goods increase.

only do we expect falling import barriers against China to generate more innovation in rich countries, but it should also occur within the same firms due to the firm-specific capital. In terms of welfare, this model suggests a benefit of lowering trade barriers against low wage countries is that it stimulates innovation, which is likely to be too low in equilibrium.[80]

In addition to predicting a within firm increase in innovation in response to a fall in import barriers against China, the trapped factor approach has two other empirical implications that we will examine. First, integration with a high wage country will have a much less positive impact on innovation. This is because imports from high-wage countries will not reduce the price of old goods relative to potential new goods as other rich countries have no comparative disadvantage in the production of new goods. In our data imports from other high wage countries do not appear to stimulate innovation, consistent with the model. A second implication of the model is that all else equal firms who have more trapped factors will respond more positively to the China shock. Using different proxies for such trapped factors (e.g. our estimates of product-specific human capital, see Appendix 4.A) we also find support for this second prediction. There may be other theories that can also rationalize the results so we do not want to over-claim for our simple model. Nevertheless, we believe it may capture some features of the stylized facts in our data and the prior case studies[81].

---

[80] In standard growth models this arises because of both knowledge externalities and the distortions induced by R&D being produced by monopolistically competitive firms. Of course, a first best solution would be to directly subsidize R&D, but in the absence of such a policy, increased trade may be a second best solution. In the model underinvestment occurs because the differentiated good sector is produced under monopolistic competition.

[81] The idea of falling opportunity costs stimulating innovation has parallels to some theories of business cycles that suggest that "bad times" can generate greater productivity enhancing activities (e.g. Aghion and Saint-Paul, 1998, or Barlevy, 2007).

4.22 Alternative Innovation models

There are several alternative models of how reducing trade barriers against low wage country goods could induce Northern innovation. First, lowering import barriers in general increases competition and competitive intensity can increase innovation. However, the effects of competition on innovation are theoretically ambiguous in general. Competition may foster innovation because of reduced agency costs (e.g. Schmidt, 1997), “higher stakes” (Raith, 2002) or lower cannibalization of existing profits.[82] However, there is a fundamental Schumpeterian force that competition lowers price-cost margins, thereby reducing the quasi-rents from innovation. We will examine competition emanating from high wage countries and show that the main effect we identify is through low wage country competition, consistent with the trapped factor model.

A second class of models stresses the importance of trade in increasing market size that will generally foster innovation incentives (e.g. Schmookler, 1966; Krugman, 1980; Acemoglu, 2008). Lower trade costs generate a larger market size over which to spread the fixed costs of investing in new technologies[83]. We will investigate these effects by examining whether European firms’ *exports to China* are associated with changes in innovation activity and show that this is not driving the imports effect we identify.

---

[82] This is the Arrow (1962) “displacement effect”. It shows up in different guises in Grossman and Helpman (1992), Aghion et al (2005)'s “escape competition” effect and the “switchover costs” of Holmes et al (2008).

[83] Recent work by Lileeva and Trefler (2010) has shown market size effects on Canadian firms of joining NAFTA.

Finally, imports could enhance innovation by enabling domestic firms to access better overseas' knowledge (e.g. Coe and Helpman, 1995 or Acharya and Keller, 2008). This may occur through the imports of intermediate inputs and supply networks (e.g. Goldberg, Khandelwal, Pavcnik and Topalova, 2010a,b)[84]. These mechanisms do not seem appropriate in the Chinese context however, as European firms have (currently) a large technological lead over China[85].

4.23 Compositional models

Perhaps an even simpler approach is to consider a framework where we keep the menu of products fixed in the economy. When trade barriers fall between the EU/US and China, the high-tech industries will grow relatively faster than low-tech industries in the EU/US. The opposite will occur in China. On empirical grounds, this simple framework is unsatisfactory, as most of the aggregate changes we observe following trade liberalization have occurred *within* rather than *between* industries. This could be explained, however, by firms operating in more finely disaggregated industries and we will show that there are strong reallocation effects whereby low-tech firms tend to shrink and exit because of China. Bernard, Jensen and Schott

[84] A related literature typically finds that productivity rises when exporting increases (e.g. Verhoogen, 2008).

[85] Eaton and Kortum (1999, 2001 and 2002) combine competition, market size and learning in a quantifiable general equilibrium trade model. For example, in Eaton and Kortum (2001) a fall in trade costs increases effective market size (which encourages innovation) but also increases competition (which discourages innovation). In their baseline model, these two forces precisely offset each other so the net effect of trade on innovation is zero. Although the Eaton-Kortum framework is powerful, it does not deal easily with one of our key results: that there is a strong effect on innovation for incumbent firms in the same sector where trade barriers fell.

(2006) show a similar result for US plants using indirectly proxies for technologies such as capital intensity.

But we also report that China induces faster technical change *within firms* and *plants,* a finding that goes beyond the existing results. In principle, firm TFP increases could be accounted for by two factors: changes in a firm's product portfolio or offshoring. First, on product switching, Bernard, Redding and Schott (2010) investigate the impact of trade liberalization in heterogeneous multi-product firms. In the face of falling trade costs with a low wage country like China, Northern firms shift their product mix towards more high-tech products (see Bernard, Redding and Schott, 2007). We will investigate this mechanism by examining how plants change their product classes, and find some evidence for this. Second, a fall in trade costs with China will mean that producers of goods that can use Chinese intermediate inputs will benefit. For example, firms may slice up the production process and offshore the low-TFP tasks to China (see for example Grossman and Rossi-Hansberg, 2008). This will have a compositional effect if the remaining activities in the home country are more technologically advanced. To investigate this mechanism we will look explicitly at offshoring to China using a method introduced by Feenstra and Hansen (1999).

Although we will show evidence that both product switching and offshoring are important in our data, neither can fully explain our core findings. In particular, about half of the China-induced increase in innovation comes from expanding the volume of patents within firms. This implies that changing composition can only be part of the story.

4.3 Data

We combine a number of rich datasets on technical change (see Appendix 4.B). Our base dataset is Bureau Van Dijk's (BVD) Amadeus that contains close to the population of public and private firms in 12 European countries. Firms in Amadeus have a list of primary and secondary four-digit industries which we use to match in the industry level trade data (the average firm had 2 primary codes, but some had as many as 10 primary and 11 secondary codes). In our main results we use a weighted average of Chinese imports across all industries that the firm operates in, but we also present robust results where we allocate the entire firm's output to a single industry.

4.31 Patents

We combined Amadeus with the population of patents from the European Patent Office (EPO) through matching by name. Patent counts have heterogeneous values so we also use future citations to control for patent quality in some specifications. We consider both a main sample of “patenters” – Amadeus firms filing at least one EPO patent since 1978 – and a wider sample where we assume that the firms unmatched to the EPO actually had zero patents.

4.32 Productivity and Research and Development (R&D)

Amadeus contains accounting information on employment, capital, materials, wage bills and sales. We calculate TFP using firms in France, Italy, Spain and Sweden because of their near population firm coverage and inclusion of materials which is needed to estimate “four-factor” TFP (materials is not a mandatory accounting item in other countries), although the results are similar using the data for all 12 countries. We estimate TFP in a number of ways, but our core method is to use a version of the Olley Pakes (1996) method applied by de Loecker (2011) to allow for trade and imperfect competition. In a first stage, we estimate production functions separately by industry across approximately 1.4 million observations to recover the parameters on the factor inputs.[86] We then estimate TFP and, in the second stage regression relate this to changes in the trade environment. As a robustness test we also allowed the production function coefficients to be different by country and industry as well as estimated at a finer level of industry aggregation which show similar results. Details of this procedure are contained in Appendix 4.C. R&D data comes from BVD's Osiris database that provides data on publicly listed firm in Europe, covering around 4,000 manufacturing firms. Of these, 459 firms report performing R&D for 5 years or more so can be used for some of the regressions.

4.33 Information technology

Harte Hanks (HH) is a multinational company that collects IT data to sell to large IT firms (e.g. IBM, Cisco and Dell). Their data is collected for roughly 160,000 establishments across 20 European countries. HH surveys establishments annually on a rolling basis which means it provides a "snapshot" of the IT stock. The data contain detailed hardware and software information. We focus on using computers per worker (PCs plus laptops) as our main measure of IT intensity because this: (i) is a physical quantity measure which is recorded in a consistent way across sites, time and countries, and (ii) avoids the use of IT price deflators which are not harmonized across countries. In robustness tests we also use alternative measures of IT such as Enterprise Resource Planning software, Groupware and Database software.

The fact that HH sells this data on to firms who use this for sales and marketing exerts a strong discipline on the data quality, as errors would be quickly picked up by clients in their sales calls. HH samples all firms with over 100 employees in each country. Thus, we do

---

[86] The number of observations in the second stage is smaller than 1.4 million because we are estimating in five-year differences.

lose smaller firms, but since we focus on manufacturing the majority of employees are in these larger firms, and we find no evidence this sampling rule biases our results.[87]

4.34 Management Practices

The management data was collected in multiple telephone survey waves between 2002 and 2010 from the London School of Economics using a team of 126 MBA-type student interviewers. Firms were interviewed for 45 minutes on average, with this interview targeted at the plant manager at the largest establishment within each firm. Management practices were scored using the Bloom and Van Reenen (2007) methodology, which scores firms on a 1 to 5 scale across 18 questions. These 18 questions span three key areas of management: *monitoring* (do firms continuously collect, analyze and use information), *targets* (do firms have balanced, well-understood and binding targets) and *incentives* (do firms reward high-performers and retrain and/or sanction poor performers). This scoring grid was developed by a leading international consulting firm based on the practices of Lean Manufacturing, the management system developed by Toyota in the 1960s and 1970s which is now becoming widely adopted across Europe, the US and Asia.

Firms were randomly sampled from the population of all manufacturing firms with 100 to 5000 employees, with a sample response rate of 44% on average (see Bloom and Van Reenen, 2010). In each wave we also resurveyed all firms from earlier survey waves to help build a management panel. In this paper we used all 579 European firms with repeat survey data.

[87] We find no systematic differences in results between firms with 100 to 250 employees and those above 250 employees, suggesting the selection on firms with over 100 employees is unlikely to cause a major bias. We also find no differences in our patenting results – where we have the full population of firms – between firms with less than and more than 100 employees. It is also worth noting that in the countries we study firms with over 100 employees account for over 80% of total employment in manufacturing.

4.35 UN Comtrade data

The trade information we use is sourced from the UN Comtrade data system. This is an international database of six-digit product level information on all bilateral imports and exports between any given pairs of countries. We aggregate from six-digit product level to four-digit US SIC industry level using the Pierce and Schott (2010) concordance. For firms that operate across multiple four digit industries we use a weighted average of imports across all sectors a firm produces in (see Appendix B)[88].

We use the value of imports originating from China (M^{China}) as a share of total world imports (M^{World}) in a country by four-digit industry cell as our key measure of exposure to Chinese trade, following the "value share" approach outlined by Bernard, Jensen and Schott (2002, 2006); i.e. we use $IMP^{CH} = M^{China} / M^{World}$. As two alternative measures we also construct Chinese import penetration by normalizing Chinese imports either on domestic production (M^{China} / D) or on apparent consumption (domestic production less exports plus imports), M^{China} / C. For domestic production we use Eurostat's Prodcom database. Compared to Comtrade, Prodcom has no data prior to 1996, so this restricts the sample period. An additional problem is that some of the underlying six-digit product data is missing (for confidentiality reasons as the industry-country cells are too small), so some missing values for domestic production had to be imputed from export data. Although we obtain similar results

88 The results are similar when we allocate firms to a single primary sector (compare Tables 1 and 5B, for example).

with all measures (see Tables 7 and A7) we prefer the normalization on world imports which does not have these data restrictions.

4.36 Descriptive Statistics

The rise of China's share of all imports to the US and the 12 European countries in our sample is remarkable. In 2000 only 5.7% of imports originated in China, but by 2007 this had more than doubled to 12.4%. This increase also varies widely across sectors, rising most rapidly in industries like toys, furniture and footwear (see Table 4.A1). Some basic descriptive statistics are shown in Table 4.A2. With the exception of the survival and worst-case bounds analyses, the regression samples condition on non-missing values of our key variables over a five year period. The exact number of observations (and average firm size) differs between samples. In the sample of firms who have patented at least once since 1978 the mean number of patents per year is one and median employment is 100. When we use the entire sample of firms with accounting data the mean number of patents falls to 0.019 and median employment to 17. R&D reporting firms are the largest of all sub-samples with 2,054 employees at the median with an average R&D intensive of 15% (recall these are all publicly listed firms whereas the other samples also include private firms). For plants with IT data, median employment is 140 and the average IT intensity is 0.58 computers per worker.

4.4 Empirical Modelling Strategy

Our empirical models analyze both the *within* firm margin of technological upgrading and the *between* firm margin of upgrading through selection effects. To investigate these we examine five broad indicators of "technology" – IT, patents, R&D, TFP and management practices.

4.41 Technical change within surviving plants and firms

Consider a basic firm-level equation for the level of technology (*TECH*) in firm *i* in industry *j* in country *k* at time *t* as:

$$\ln TECH_{ijkt} = \alpha IMP^{CH}_{jkt-l} + \beta x_{ijkt} + \varepsilon_{ijkt} \tag{4.1}$$

TECH will be interpreted broadly and measured using a number of indicators such as patented innovations[89], R&D spending, IT, TFP and management practices. We measure IMP^{CH}_{jkt} mainly as the proportion of imports (*M*) in industry *j* and country *k* that originate from China $(M^{China}_{jk} / M^{World}_{jk})$ and the x_{ijkt} are a set of control variables such as country dummies interacted with time dummies to absorb macro-economic shocks. The trade-induced technical change hypothesis is that $\alpha > 0$. Note that we allow for a dynamic response in equation (1) depending on the lag length indicator *l*. Our baseline results will use $l = 0$ to be consistent with the other technology equations, but we show the differences in results to alternative lag lengths in sub-section 4.5[90]

Since there may be many unobservables that are correlated with the firm (and industry's) level of technology and imports that different across firms but broadly constant over time, we will control for these by including a fixed effect and estimate:

[89] Because of the zeros in patents when taking logarithms we use the transformation *PATENTS* = 1 + *PAT* where *PAT* is the count of patents. The addition of unity is arbitrary, but equal to the sample mean of patents. We also compare the results with fixed effect Negative Binomial count data models below which generated similar results (see Table 6).

[90] For patents, the largest effects appear after three years (see Table A6) which is consistent with the idea that most firms take a few years to obtain innovations from their increased R&D spending.

$$\Delta \ln TECH_{ijkt} = \alpha \Delta IMP^{CH}_{jkt-l} + \beta \Delta x_{ijkt} + \upsilon_{ijkt} \tag{4.2}$$

We use Δ to denote the long (usually five year) difference operator. Rapid growth in the Chinese import share is therefore used as a proxy for a rapid increase in trade competition from low wage countries. The growth of Chinese imports may still be related to unobserved shocks, υ_{ijkt} so we consider instrumental variables such as the removal of quotas when China joined the WTO to evaluate potential endogeneity biases. We maximize the use of the data by using overlapping five-year differences (e.g. 2005-2000 and 2004-1999) but since we cluster at the country-industry pair level (or sometimes just industry level) this is innocuous. We report some results using non-overlapping five-year differences and specifications in levels (e.g. fixed effect Negative Binomial models).

4.42 Technological upgrading through reallocation between plants and firms

In addition to examining whether Chinese import competition causes technological upgrading *within* firms we also examine whether trade affects innovation by reallocating economic activity *between* firms by examining employment and survival equations. As discussed in the Section III, compositional models would predict that China would cause low-tech plants to shrink and die, as they are competing most closely with Chinese imports. Consequently, we estimate firm employment growth equations of the form:

$$\Delta \ln N_{ijkt} = \alpha^{N} \Delta IMP^{CH}_{jkt} + \beta^{N} \Delta x^{N}_{ijkt} + \gamma^{N} (TECH_{ijkt-5} * \Delta IMP^{CH}_{jkt}) + \delta^{N} TECH_{ijkt-5} + \upsilon^{N}_{ijkt}$$
(4.3)

where the coefficient α^{N} reflects the association of jobs growth with the change in Chinese imports, which we would expect to be negative (i.e. α^{N} <0) and *TECH* is the relevant technology variable (e.g. patenting). We are particularly interested in whether Chinese import competition has a larger effect on low-tech firms, so to capture this we include the interaction of ΔIMP_{jkt}^{CH} with the (lagged) technology variables. If Chinese trade has a disproportionately negative effect on low-tech firms we would expect $\gamma^{N} > 0$.

Equations (4.2) and (4.3) are estimated on surviving firms. However, one of the effects of Chinese trade may be to reduce the probability of plant survival. Consequently, we also estimate:

$$SURVIVAL_{ijkt} = \alpha^{S}\Delta IMP_{jkt}^{CH} + \beta^{S}\Delta x_{ijkt}^{S} + \gamma^{S}(TECH_{ijkt-5} * \Delta IMP_{jkt}^{CH}) + \delta^{S} TECH_{ijkt-5} + \upsilon_{ijkt}^{S}$$

(4.4)

which is defined on a cohort of firms (or establishments) who were alive in a base period and followed over the next five years. If these establishments (or firms) survived over the subsequent five years we define $SURVIVAL_{ijkt}$ = 1 and zero otherwise. If Chinese imports do reduce survival probabilities, we expect $\alpha^{S} < 0$ and if high-tech plants are more protected we expect $\gamma^{S} > 0$.

To complete the analysis of between firm effects we would also need an entry equation. The fundamental problem is that there is no "initial" technology level for entering firms. We cannot use the current observed technology level ($TECH_{ijkt}$) as this is clearly endogenous (in equations (4.3) and (4.4) we use lagged technology variables under the assumption that technology is weakly exogenous). We can address the issue of entry indirectly, however, by estimating an industry-level version of equation (2):

$$\Delta TECH_{jkt} = \alpha^{IND}\Delta IMP_{jkt}^{CH} + \beta^{IND}\Delta x_{jkt} + \upsilon_{jkt}^{IND} \qquad (4.5)$$

where the coefficient on Chinese imports, α, in equation (4.5) reflects the combination of within firm effects from equations (4.1) and (4.2), the reallocation effects from equations (4.3) and (4.4), and the unmodelled entry effects. In examining the magnitude of the Chinese trade effects, we will simulate the proportion of aggregate technical change that can be accounted for by Chinese imports using equations (4.2)-(4.4) and break this down into within and between components. We will also compare the micro and industry estimates of equation (4.5) which give an alternative estimate of the within and between effects, including entry.

4.5 Results

4.51 Within firm and within plant results

Table 4.1 presents our core results: within firm and within plant measures of technical change. All columns control for fixed effects by estimating in long-differences and country-specific macro shocks by including a full set of country dummies interacted with a full set of time dummies. Our key measure of innovation, patents, is the dependent variable in column (1). The coefficient suggests that a 10 percentage point increase in Chinese import penetration is associated with a 3.2% increase in patenting. Since jobs fell in those industries affected by Chinese imports (see Table 4.3) we underestimate the growth in patent intensity (patents per worker) by not controlling for (endogenous) employment. If we also include the growth of employment in column (1), the coefficient (standard error) on imports is slightly larger at 0.387 (0.134).[91] Note that our pooling across multiple overlapping years to construct five-year

[91] The coefficient (standard error) on employment in the patents equation was 0.015(0.008) implying that larger firms have a higher volume of patents. If we include the ln(capital/sales) ratio as well as ln(employment) in the regression this barely shifts the results (the coefficient on Chinese imports is 0.370 with a standard error of 0.125). Thus, the correlation with Chinese trade is not simply an increase in all types of capital, but seems related specifically to technical change. The other results in the table are all

differences is largely innocuous as we are clustering the standard errors by country-industry pair. For example if we use only the last five year difference the qualitative results are similar. In this experiment the coefficient (standard error) is 0.591(0.201) for patents; 0.314(0.077) for IT; and 0.400 (0.079) for TFP.

A concern with patenting as an innovation indicator is that firms may simply be taking out more patents to protect their existing knowledge in the face of greater Chinese competition. To test this "lawyer effect" we also look at citations per patent – if firms are now patenting more incremental knowledge for fear of being copied by the Chinese, the average quality of their patents should fall, so citations per patent should drop. The results on citations per patents in Table 4.A3 show, in fact, that Chinese competition does not lead to a fall in citations. The coefficient on Chinese imports is actually positive (but insignificant).

In column (2) of Table 4.1 we examine IT intensity and find a positive and significant coefficient on Chinese imports. We use computers per employee as our main measure of IT diffusion as this is a good indicator of a general-purpose technology used widely across industries. However, in Table 4.A4 we investigate other measures of IT – the adoption of Enterprise Resource Planning, database software, and groupware tools – and also find positive coefficients on Chinese imports.

Column (3) of Table 4.1 uses R&D as the outcome and also shows a large and significant increase in firm-level R&D expenditure when Chinese imports rise, which is more evidence that the increase in innovation observed in column (1) is not due to firms merely taking

---

robust to controlling for employment growth – see Table 4.5 below for more analysis controlling for firm and industry size.

out more intellectual property protection. Column (4) uses a wider measure of technical change as the dependent variable, TFP growth, and again establishes a positive and significant association with Chinese imports growth. The final column delves uses the latest version of the management practices data first described by Bloom and Van Reenen (2007).[92] Hence, Chinese trade competition also appears to stimulate the rapid adoption of modern (Lean) management practices.

As we discuss in Section 4.6 below the magnitudes are economically significant: a 10 percentage point increase in Chinese imports is associated with a 3.2% increase in patenting, a 3.6% increase in IT, a 12% increase in R&D, a 2.6% increase in TFP and a 0.8 (1.38 standard deviation) increase in the management index. Note that the latter is an extremely large change given that that the average gap between US and Indian management practices is 0.85.

4.52 Endogeneity: the problem of unobserved technology shocks

An obvious problem with estimating these equations is the potential endogeneity of Chinese imports due to unobserved technology shocks correlated with the growth of Chinese imports. To address this we consider three alternative strategies to control for these unobserved shocks: (i) using the natural experiment of China joining the WTO, (ii) constructing an IV from initial conditions and (iii) controlling for industry time trends. The smaller size of the datasets for R&D and management makes it infeasible to implement these identification strategies so we focus mainly from this point forwards on the three large-sample technology measures: patents, IT and TFP.

---

[92] We have up to three panel data observations per firm between 2010 and 2002 across the European countries considered here (see Appendix B) so can only use shorter (three year) differences.

China joining the WTO as a quasi-experiment - One identification strategy is to use the accession of China to the WTO in 2001, which led to the abolition of import quotas on textiles and apparel. We discuss this in detail in Appendix D, but sketch the idea here. The origin of these quotas dates back to the 1950s when Britain and the US introduced quotas in response to import competition from India and Japan. Over time, this quota system was expanded to take in most developing countries, and was eventually formalized into the Multi-Fiber Agreement (MFA) in 1974. The MFA was itself integrated into GATT in the 1994 Uruguay round, and when China joined the WTO in December 2001 these quotas were eliminated in two waves in 2002 and 2005 (see Brambilla, Khandelwal and Schott, 2010). Since these quotas were built up from the 1950s, and their phased abolition negotiated in the late 1980s in preparation for the Uruguay Round, it seems plausible to believe their level in 2000 was exogenous with respect to future technology shocks. The level of quotas also varied quasi-randomly across four-digit industries[93] – for example, they covered 77% of cotton fabric products (SIC 2211) but only 2% of wool fabric products (2231), and covered 100% of women's dresses (2334) but only 5% of men's trousers (2325). This variation presumably reflected the historic bargaining power of the various industries in the richer countries in the 1950s and 1960s when these quotas were introduced, but are now uncorrelated to any technology trends in the industries as we show below.

When these quotas were abolished this generated a 240% increase in Chinese imports on average within the affected product groups. In fact, this increase in textile and

[93] The quotas were actually imposed at the six-digit level that we aggregated up to the four-digit industry level weighting by their share of world imports calculated in the year 2000 (the year before WTO accession).

apparel imports was so large it led the European Union to re-introduce some limited quotas after 2005.[94] Since this re-introduction was endogenous, we use the initial level of quotas in 2000 as our instrument to avoid using the potentially endogenous post-2005 quota levels. Although the quota-covered industries are considered low-tech sectors, European firms in these industries generated 21,638 patents in our sample. In Appendix 4.D we give several examples of such patents taken out by European firms.

Panel A of Table 4.2 uses this identification strategy of China's accession to the WTO.[95] Since this is only relevant for textiles and clothing, we first present the OLS results for these sectors for all the technology indicators in columns (1), (4) and (7). In column (1) there is a large positive and significant coefficient on the Chinese trade variable, reflecting the greater importance of low wage country trade in this sector. Column (2) presents the first stage using the (value-weighted) proportion of products covered by quotas in 2000. Quota removal appears to be positively and significantly related to the future growth of Chinese imports. Column (3) presents the IV results that show a significant effect of Chinese imports on patents with a higher coefficient than OLS (1.86 compared to 1.16).

---

[94] The surge in Chinese imports led to strikes by dock workers in Southern Europe in sympathy with unions from the apparel industry. The Southern European countries with their large apparel sectors lobbied the European Union to reintroduce these quotas, while the Northern European countries with their larger retail industries fought to keep the quota abolition. Eventually temporary limited quotas were introduced as a compromise, which illustrates how the abolition of these quotas was ex ante uncertain, making it harder to pick up anticipation effects.

[95] In Panels A and B of Table 4.2 we cluster by four-digit industry as the instruments have no country-specific variation. We also drop years after 2005 so the latest long difference (2005-2000) covers the years before and after China joined the WTO.

Columns (4)-(6) repeats the specification but uses IT intensity instead of patents as the dependent variable. Column (4) shows that the OLS results for IT are also strong in this sector and column (5) reports that the instrument has power in the first stage. The IV results in column (6) also indicate that the OLS coefficient appeared downward biased.[96] The final three columns repeat the specification for TFP showing similar results to patents and IT. So overall there is a large OLS coefficient for patents, IT and TFP, but an even larger IV coefficient and certainly no evidence of upward bias for OLS.[97]

There are several issues with the specification. First, the regressions all use the actual flow of Chinese imports to reflect the threat of import competition. However, an advantage of the IV estimates is that by replacing the actual flow of imports by the predicted flow based on quota relaxation, this more accurately reflects the *threat* of Chinese competition.[98] Second, one could argue that firms will be adjusting their innovation efforts earlier in response to *anticipation* of quota relaxation. However, at the time there was considerable uncertainty over whether the liberalization would actually take place. A common view was that even if there was an abolition of quotas this would be temporary, as to some extent it was with the temporary

---

[96] If we repeat the IV specification of column (6) but also condition on employment growth the coefficient on Chinese imports is 0.687 with a standard error of 0.373. Dropping all the four-digit sectors that had a zero quota in 2000 uses only the continuous variation in quotas among the affected industries to identify the Chinese import effect. Although this regression sample has only 766 observations, this produces a coefficient (standard error) under the IV specification of 2.688(1.400) compared to an OLS estimate of 1.238(0.245).

[97] The Hausman tests fail to reject the null of the exogeneity of Chinese imports for the patents and IT equations, but does reject for the TFP equation (p-values of 0.342, 0.155 and 0.001 respectively).

[98] In the reduced forms the coefficient (standard error) on Chinese imports was 0.201(0.091), 0.163(0.038) and 0.129(0.018) in the patents, IT and TFP equations. Regressions include country dummies times year dummies.

reintroduction of some quotas in 2006. We discuss this issue in more detail in Appendix D where we show that there is no significant correlation of the quota instrument with technical change or Chinese imports *prior* to the 2001 WTO accession.[99] This placebo experiment also addresses the concern that quota intensity is proxying some other trend correlated with Chinese import growth.

To further examine this issue we include lagged Chinese import growth (1995-2000) as an additional control in Table 4.2. The coefficients are robust to this.[100] The most rigorous test is to include lags of both technology and Chinese imports in the regression, which we do in Table 4.A5. We use the TFP specifications as we have the largest time series of data in order to condition on the pre-policy variables. Column (1) of Table 4.A5 repeats the specification from the final column of Table 4.2 Panel A. Column (2) conditions on the balanced panel where we observe firms for 10 years and shows that the results are robust even though we have only two-thirds of the industries. Column (3) includes the two pre-policy variables, the lagged growth of imports and the lagged growth of TFP. The coefficient on lagged imports is insignificant, but lagged TFP is negative and significant. Importantly, the coefficient on current Chinese import growth remains positive and significant, actually rising from 1.49 to 1.61. The negative coefficient on the lagged dependent variable is expected due to mean reversion, so we

---

[99] For example, to test for anticipation effects we regressed the level of the imports quota on the growth of technology 1996-2000 *prior* to WTO accession in 2001. All the coefficients on technology were small and insignificant suggesting no anticipation effects. The coefficient (standard error) was 0.096(0.177) for patents and 0.024(0.031) for TFP. We do not have IT data before 2000 so cannot implement this placebo test.

[100] For example in column (6) the coefficient on lagged imports is positive (0.168) but insignificant and the coefficient on Chinese import growth remains positive and significant (1.792 with a standard error of 0.421).

also report the results of instrumenting this with the firm's 1996 level of TFP. This reverses the sign of the coefficient on the lag, suggesting a positive relationship between past and present TFP. Again the coefficient on Chinese imports is essentially unchanged, suggesting that the significant impact of imports instrumented by WTO quota abolition is not proxying for pre-existing industry trends.

Initial conditions as instrumental variables - A disadvantage of the quota-based instrument is that we can only construct the instrument for the affected industries (textiles and clothing), so we consider a second identification strategy. The overall increase in Chinese imports is driven by the exogenous liberalization being pursued by Chinese policy makers. The industries where China exports grew more depended on whether the industry is one in which China had a comparative advantage. For example, if we consider the growth of Chinese imports in Europe between 2000 and 2005, sectors in which China was already exporting strongly in 1999 are likely to be those where China had a comparative advantage – such as textiles, furniture and toys – and are also the sectors which experienced much more rapid increase in import penetration in the subsequent years (see Table 4.A1). Consequently, high exposure to Chinese imports in 1999 can be used (interacted with the exogenous overall growth of Chinese imports, ΔM^{China}) as a potential instrument for subsequent Chinese import growth. In other words we use ($IMP^{CH}_{jt-6} * \Delta M^{China}_{t}$) as an instrument for ΔIMP^{CH}_{jkt} where IMP^{CH}_{jt-6} is the Chinese import share in industry *j* in the EU and US. Note that we do not make IMP^{CH}_{jt-6} specific to country *k* to mitigate some of the potential endogeneity problems with initial conditions.[101]

[101] This identification strategy is similar to the use of "ethnic enclaves" by papers such as Card (2001) who use the proportion of current immigrants in an area as an instrument for future immigrants.

A priori, the instrument has credibility. Amiti and Freund (2010) show that over the 1997 to 2005 period at least three quarters of the aggregate growth of Chinese imports was from the expansion of existing products rather than from adding new products. Similarly, Brambilla et al (2010) find this was true when focusing on textiles and clothing after 2001.[102]

Column (1) of Table 4.2 panel B re-presents the basic OLS results for patents. Column (2) presents the first stage for the instrumental variable regressions. The instrument is strongly correlated with the endogenous variable, the growth of Chinese import intensity. Column (3) presents the second stage: the coefficient on Chinese imports is 0.495 and significant.[103] Columns (4) through (6) repeat the experiment for IT. In column (6) the coefficient on Chinese imports is positive and significant and above the OLS estimate. In the final column (9) for TFP, the IV coefficient is again above the OLS estimate.[104] Taking Panels A and B of Table 4.2 as a whole, there is no evidence that we are under-estimating the effects of China on technical change in the OLS estimates in Table 1. If anything, we may be too conservative.[105]

---

[102] This appears to be common in several countries- e.g. Mexico after NAFTA (e.g. Iacovone and Javorcik, 2008).

[103] Unsurprisingly the results are more precise if we combined the initial conditions and quota instruments together. For example in column (3) the coefficient (standard error) on patents is 2.322 (0.990). Furthermore, we cannot reject the null that the instruments are valid using a Hansen over-identification test. The p-values for rejection of instrument validity are 0.438 for the patent equation, 0.330 for the IT equation and 0.948 for the TFP equation.

[104] If we use the initial conditions estimator for R&D following the column (9) specification we find a point estimate (standard error) of 1.179 (0.582).

[105] The downward bias on OLS of trade variables is also found in Auer and Fisher (2010) who examine the impact of trade with less developed countries on prices. They use a variant of an initial conditions

Controlling for technology shocks using industry trends. A third way to control for unobservable technology shocks is to include industry trends. We do this in Panel C of Table 4.2 by including a set of three-digit industry dummies in the growth specifications. In column (1) we reproduce the baseline specification for patents and column (2) includes the industry trends. We repeat this for each of the technology variables. Although the magnitude of the coefficient on Chinese imports is smaller in all cases, it remains significant at the 10% level or greater across all three specifications. Note that the industry trends are jointly insignificant in all three cases. It is unsurprising that the coefficient falls as we are effectively switching off' much of the useful variation and exacerbating any attenuation bias. Furthermore, although including these industry dummies may deal with omitted trends, they do not deal with reverse causality, which, as argued above will cause a downward bias on the coefficient of interest. The WTO instrument is superior in this respect as in principle it deals with both reverse causality and omitted variables.

The results are generally robust to even tougher tests. If we include four digit industry trends the coefficient (standard errors) in the patent, IT and TFP regressions are 0.185(0.125), 0.170(0.082) and 0.232(0.064). If we include three digit dummies interacted with country dummies the results are: 0.274(0.101); 0.176(0.08) and 0.167(0.052). Hence, the primary source of identification is (i) multi-product firms who face differential industry effects in addition to their primary sector and (ii) the acceleration of import growth and technology. The continued importance of the trade variable even after this tough test is remarkable.

---

estimator based on the industry's labor intensity. Like them, we also find important import effects on prices (see sub-section 4.6).

Summary on endogeneity - The main concern in interpreting the technology-trade correlation in Table 1 as causal is that there are unobserved technology shocks. The evidence from Table 4.2 is that controlling for such potential endogeneity concerns in a variety of ways does not undermine a causal interpretation of the impact of Chinese imports on technical change in the North.

4.53 Reallocation effects: jobs and survival

Table 4.3 examines reallocation effects by analyzing employment growth in Panel A and survival in Panel B. We first examine the basic associations in column (1) of Panel A, which suggests a strong negative effect of Chinese imports - a 10 percentage point increase in imports is associated with a 3.5% fall in employment. In addition, high-tech firms (as indicated by a high level of lagged patents per worker) were more likely to grow. Most importantly, the interaction of Chinese trade and lagged patent stock enters with a positive and significant coefficient in column (2). This suggests that more high-tech firms are somewhat shielded from the effects of Chinese imports. In columns (3) and (4) we repeat the estimates but for the "patenters" sample rather than all firms (i.e. those firms who had at least one patent since 1978) and find a similar result: firms with a high lagged patent stock had less job falls following a Chinese import shock.[106] In columns (5) and (6) we run similar employment estimations using the initial level of IT and TFP and again find similar positive interaction terms, suggesting high-tech firms are somewhat protected from the effects of Chinese import competition.

---

[106] Furthermore, this result is not driven by the inclusion of employment in our patent stock measure. To test this we estimated both a model where employment was removed from the denominator (that is, a simple patent stock measure) and a model that include lagged employment and its interaction with Chinese imports. The estimate of our technology-imports interaction terms for these models were 0.192(0.086) and 0.160(0.083) respectively.

We also examined the dynamic effects of Chinese imports on employment and compared this to the impact on technology. Table 4.A6 explores the timing for patents by moving from a lag-length of 5 years in column (1) to a lag-length of zero years in column (6) as in our baseline model. Chinese imports appear to have the largest impact on patents after about three years. Panel B of Table 4.A6 shows the same results for employment, where we see the largest impact for Chinese imports is contemporaneously. This is consistent with the idea that firms respond to Chinese imports by cutting employment while also initiating innovative projects. These innovation projects appear to take around three years on average to produce innovations that are sufficiently developed to be patented.

Panel B of Table 4.3 examines survival. We consider a cohort of firms and plants alive in 2000 and model the subsequent probability that they survived until 2005 as a function of the growth of industry-wide Chinese imports and the initial technology levels. Column (1) shows firms facing higher rates of Chinese import growth are less likely to survive - a ten percentage point increase in Chinese imports decreases the survival probability by 1.2 percentage points. Since the mean exit rate in our sample period is 7-percentage point this represents a 17% increase in exit rates. Column (2) analyzes the interaction term between Chinese import growth and lagged patents and finds again a positive "shielding" effect – firms with a low initial patent stock have a significantly higher change of exiting when faced by an influx of Chinese imports. In columns (3) and (4) we re-estimate these specifications using only patenting firms and again find a significant positive interaction between lagged patent stocks and Chinese imports[107]. Columns (5) and (6) shows that there are also positive interaction effects when we use IT or TFP as alternative measures of technology, although these are not significant at the 5% level. Further investigation reveals that the main effect is coming from firms in the bottom quintile of the technology distribution who were significantly more likely to exit because of Chinese import competition.[108] These findings on the impact of low wage country imports on reallocation is consistent with those found in US manufacturing establishments in Bernard, Jensen and Schott (2006) using indirect measures of technology (capital intensity and skills) for the pre-1997 period in the US.

4.6 Magnitudes: Industry-Level results, Selection and General Equilibrium Effects

Taking all these results together we have a clear empirical picture of the role of Chinese imports in increasing technological intensity both within firms (Tables 1 and 2) and between firms by reallocating output to more technologically advanced firms (Table 4.3). We now turn to the economic magnitude of these effects.

4.61 Magnitudes

We can use the regression coefficients to perform some partial equilibrium calculations to quantify how much of the aggregate change in technology China could account for and to gauge the relative importance of within and between firm effects (details in Appendix E). In summary, for patents per employee we apply the coefficients from all our regressions with the

[107] We have re-estimated all these results with the IV strategies discussed in the previous section and, as with the technology equations, all results are robust.

[108] For example, estimating column (5) but using the lowest quintile of the IT intensity distribution rather than the linear IT intensity gave a coefficient (standard error) of 0.214 (0.102) on the interaction.

empirical growth of Chinese imports to predict growth in patent intensity and then divide this by the actual growth in aggregate patent intensity in our sample. For IT and TFP we follow a similar exercise, again applying our regression coefficients to get a predicted increase from China and dividing by the total increase in aggregate data.

In Table 4.4 we see that over the 2000-2007 period Chinese imports appear to have accounted for about 14.7% of the increase in aggregate patenting per worker, 14.1% of the increase in IT intensity and 11.8% of TFP growth in European manufacturing. The predicted impact of Chinese imports appears to increase over this period. For example, we estimate that Chinese imports accounted for 13.9% of the increase in patents per employee over the 2000-04 period but 18.7% over the 2004-2007 period. The reason for this acceleration is clear in Figure 4.1, where Chinese import growth has rapidly increased over this period. Table 4.4 also shows for patents the contributions of the within and between components are roughly equal which is consistent with the literature on trade liberalization (e.g. Pavcnik, 2002). For IT and productivity, the within component is larger which may be because the adjustment costs are lower in response to the more gradual growth of Chinese imports over the 2000's compared to the "shock" trade liberalizations examined in places like Chile and Columbia.

4.62 Industry level results

In Table 4.5 we re-estimate our technology regressions at the industry level in Panel A and at the firm level in Panel B.[109] This provides another approach to comparing the within firm

[109] The firm-level results are identical to those in Table 1 for IT and R&D. The patents and TFP results differ somewhat from Table 1 because we exploited the multi-industry information at the firm level to construct a weighted average of Chinese imports in the main results. By contrast, in Table 5 we allocate a firm to its primary four-digit industry (Panel B) for comparability to the industry level results (Panel A). See Appendix B for details.

and between firm magnitudes of the impact of trade with China, since the industry level magnitudes capture both effects while the firm level magnitudes capture only the within effects. In addition to being a cross check on the magnitudes as estimated from the full set of equations, the industry-level estimates include any effect of China on entry.[110] For example, if Chinese competition discourages entry of innovative firms within an industry, then the calculations in Table 4.4 will over-estimate the impact of trade on technical change. By contrast, the industry level aggregates are the stock of firms so include all growth from entrants as well as survivors.

Table 4.5 starts by examining outcomes where we expect Chinese trade to have a negative impact: prices, employment and profitability. We use producer prices as a dependent variable in column (1) of Panel A (there is no firm-level price data) and observe that Chinese imports are associated with large falls in prices in the most affected industries, consistent with Broda and Romalis (2009). Column (2) uses employment as the dependent variable and shows a larger negative effect at the industry level (Panel A) than the firm level (Panel B) consistent with the evidence from Table 4.3 that there is a trade effect on exit probabilities.[111] Column (3) contains the results for profitability (profits before tax, interest and dividends divided by revenue) and shows that industry and firm profits have fallen significantly (the smaller firm-level coefficient is the usual selection effect due to the least profitable firms being the first to exit). This negative profitability effect is important, as it is consistent with the idea that Chinese

[110] Atkeson and Burstein (2010) stress this as one of the main problems with firm-level analysis of trade. See also Arkolakis, Costinot and Rodríguez-Clare (2010).

[111] Interestingly, Autor, Dorn and Hansen (2012) looking at US labor markets find Chinese import competition not only reduces employment and wages, but also increases transfer payments for disability and unemployment.

imports are causing an increase in competitive pressure in the industry (as assumed in the "trapped factor" model). If Chinese import share was instead only proxying some greater ability to offshore (which if properly measured it should not as these are Chinese imports in the firm's *output* market not its *input* market), then we would expect the coefficient to be positive as this should enhance rather than inhibit profitability. We discuss offshoring in more detail in sub-section VII.D below.

In columns (4) to (10) of Table 4.5 we show results for our technology measures - patents, IT, R&D and TFP. At the industry level (Panel A) we find that Chinese import competition is significantly associated with increases in all of these measures of technology. Note that these industry-level results are based on data that includes all firms and plants at a given point in time, rather than just survivors. In Panel B columns (4) to (10) confirm that the firm level results show similar strong associations between Chinese import growth and technology, but with magnitudes between one-half to two-thirds of those at the industry level, broadly consistent with the share of the within firm component shown in the Table 4.4 magnitude calculations. This suggests that any entry effects omitted from the firm-level results, but included in the industry level results, must be relatively small given the similarity of the magnitudes[112].

4.63 Dynamic Selection Bias

A concern with our finding of positive effects of Chinese imports competition on within firm technical change is that it reflects dynamic selection bias. For example, it may be that firms

---

[112] For example, the magnitude of the within industry level effects 2000-2007 for patents, IT and TFP are 12.5%, 10.8% and 16.1%, very similar to the equivalent firm-level values of 14.7%, 14.1% and 11.8% as shown in Table 4.

who know that they are technologically improving are less likely to exit in the face of the Chinese import shock. This could generate our positive coefficients in Table 1. Note that our industry-level results in Table 4.5A are robust to this problem as it examines aggregate innovation. Dynamic selection bias would mean, however, that we allocate too much of this aggregate industry effect to the within firm component and too little to the reallocation component in the calculations of Table 4.4.

Appendix F gives a formal statement of the dynamic selection problem and suggests two ways of tackling it by (i) bounding the selection bias and (ii) a control function approach. First, we can place an upper bound on the magnitude of the dynamic selection effects by exploiting the fact that the number of patents can never fall below zero. We create pseudo observations for firms who exit and give them a value of zero patents for all post exit periods until the end of the sample in 2005. This is a "worst case bounds" bounds approach (see Manski and Pepper, 2000 or Blundell et al, 2007) as the effect of trade could never be less than this lower bound.

Table 4.6 implements this method. We first report the baseline results of Table 1 column (1) and then report the results for the worst-case lower bounds in column (2). Note that as well as additional observations on our existing 8,480 firms we also obtain additional firms as we now can construct a five-year difference even for firms with less than five years of actual patenting data by given them zeros for the years after they exit. Dropping firms with less than five years of data is another possible source of selection bias that is addressed by this method.[113].

[113] A total of 658 firms some history of patenting exited to bankruptcy in our sample. 406 of these were already in the main sample of 8,480 firms and 30,277 observations (Table 1, column (1)). The additional 252 of the 658 exiting firms were outside the main sample because they reported less than five

Our results appear to be robust to these potential selection bias problems as the coefficient on Chinese imports in column (2) remains positive and significant and has fallen only by less than one-sixth, from 0.321 to 0.271.

Since patents are counts we also consider a Negative Binomial model. It is less straightforward to deal with fixed effects in such models than in our baseline long-differences models, especially with weakly exogenous variables like Chinese imports (e.g. the Hausman, Hall and Griliches, 1984, fixed effect Negative Binomial model requires strict exogeneity). We use the Blundell et al (1999) method of controlling for fixed effects through pre-sample mean scaling for the baseline model. This estimator has proven attractive in the context of patent models and exploits the long pre-sample history of patents to control for the fixed effect (we have up to 23 years of pre-sample patent data). More details of the estimation technique are in Blundell et al (2002) and the textbook by Cameron and Trevidi (2005).

Column (3) implements the Negative Binomial model and shows that the coefficient on imports is similar to the baseline results with a positive and significant coefficient that is if anything slightly higher than the long differenced results. Column (4) shows that the worst-case lower bounds are again not much lower than the baseline, with the effect falling from 0.397 to 0.389.

We conclude from Table 4.6 that the dynamic selection problem is not causing us to substantially overestimate the impact of Chinese competition causing a within firm increase in the volume of innovative activity.

---

consecutive observations so that a five-year difference in patenting could not be defined. The increase in observations from 30,277 in column (1) to 31,272 in column (2) are the additional observations on these 658 exiting firms.

This worse case bounds approach will not work for TFP as it does not have a lower bound of zero. However, the approach that we have taken to calculate TFP already includes a control function approach to remove the bias associated with selection in production functions following Olley and Pakes (1996). Details of this are in Appendix C. Thus, our TFP estimates should be robust to this selection problem.

4.64 General Equilibrium and Welfare

We cannot prematurely jump to aggregate welfare conclusions from the results in the paper. Atkeson and Burstein (2010) argue that lowering trade costs may lead to the exit of low productivity domestic firms, but it will deter product innovation through new entry. In Ossa and Hsieh (2010) the reduction of barriers to Chinese imports raises average European firm productivity (as we find), but lowers the average quality of Chinese exporters to the EU. While Arkolakis et al (2008, 2010) argue that the standard gains to trade summarized in the ratio of exports to GDP are not fundamentally altered in a wide class of models that allow for heterogeneous firms. More subtly, the innovation response in rich countries in sectors where China has comparative advantage (like textiles), might reduce the standard Ricardian gains from trade (Levchenko and Zhang, 2010).[114]

[114] Another argument is that increased innovation from Chinese trade drives up the wages of R&D scientists leading to no net increase in innovation, We believe this fully offsetting increase in R&D prices is unlikely. First, much of the improvements we identify do not require large increases in R&D scientists – the incremental changes in IT, TFP, management practices and patenting may require more skilled workers, but not more scientists. Second, it is unlikely the supply curve of R&D scientists is completely vertical – workers for innovation-related tasks can be imported from overseas and redeployed from other activities. Bloom, Griffith and Van Reenen (2002), for example, showed that the number of R&D employees rose in countries that introduced fiscal incentives for R&D even in the short-run.

Our empirical models are partial equilibrium and do not capture all of the complex welfare effects of trade with China. Therefore, what they directly estimate is the impact of increasing trade on innovation on an industry-by-industry basis. This is directly relevant for typical trade policy question, such as the costs of putting quotas on imports in any particular industry.

Nevertheless, we think that our results are also suggestive of a positive aggregate effect of Chinese trade on innovation as implied by standard endogenous growth models (such as the trapped factor model of Bloom et al, 2012, building on Romer and Riviera-Batiz, 1992)[115]. First, the within firm effects of Chinese imports on innovation is at least as large as the between firm-reallocation effects (see Table 4.4). Second, any possible falls in innovation through lower entry within the same industry cannot offset our main results. For example, the net effect of China (including all entry effects) is on average positive in the industry-level analysis of Table 4.5 Panel B.

4.7 Extensions and Robustness

In Section II we discussed several models of trade induced technical change. The trapped factor model, amongst others, suggested that innovation should rise when faced by greater import competition and should occur for firms facing the largest trade shock. The trapped factor model also implies that the innovation response (i) should be weaker for import competition from high wage countries, (ii) larger for firms more subject to the trapped factor problem. We investigate these further implications in the next two sub-sections, examine skills

[115] It may be that there is "too much" innovation of course, so slowing down innovation has positive welfare effects. However, most empirical estimates have found that there is a socially sub-optimal level of R&D and innovation (e.g. Jones and Williams, 1998).

as another outcome in sub-section C and finally examine three alternative theories of Section II relating to offshoring, product switching and export-led innovation in sub-sections VII.D to VII.F.

4.71 Low wage vs. high wage country trade

Our key measure of Chinese import competition is the share of total imports originating in China. An alternative approach is to normalize Chinese imports by a measure of domestic activity such as production or apparent consumption. These alternative normalizations are presented in Table 4.A7. Although the magnitude of the coefficients changes as the mean of the imports variable is different, the qualitative and quantitative results are remarkably similar.[116]

Using these alternative definitions of Chinese imports also allows us to separately distinguish the impact of Chinese imports from all other low wage country imports and high wage country imports. We define low wage countries as those countries with GDP per capita less than 5% of that in the US between 1972 and 2001. On this definition, the increase in non-Chinese low wage imports (as a proportion of all imports) 1996-2007 was close to zero (0.005), whereas China's growth was substantial (see Figure 4.1).

Table 4.7 presents some analysis of using measures of Chinese imports normalized by domestic production. The dependent variable is the change in patents in Panel A, the change in IT in Panel B and the change in TFP in Panel C. Column (1) simply shows what we have already seen – Chinese import penetration is associated with significantly greater technical change. Column (2) includes the non-Chinese low wage country import penetration

---

[116] For example, a one standard deviation increase in the import share in Table 1 column (1) is associated with a 10% increase in patenting. By contrast, a one standard deviation increase in the import share in column (1) of Panel B in Table A7 is associated with a 14% increase in patenting.

measure. The coefficient is insignificantly different from the Chinese imports coefficient in all panels. When we include all low wage country import penetration instead of just China in column (3) we obtain similar coefficients to those in column (1), with a positive and significant coefficient for all three technology measures. We conclude that China is qualitatively no different from other low wage countries - it is just the largest trade shock from low wage countries in recent decades.

Column (4) of Table 4.7 includes the growth of imports from high wage countries. The coefficient is positive in all panels, but insignificant. High wage imports are also easily dominated by Chinese imports when both are included in column (5). Column (6) uses total import penetration that is positive but again dominated by China in column (7). One concern is that the endogeneity bias may be greater for high wage country imports than Chinese imports. We followed Bertrand (2004) and used trade-weighted exchange rates as an instrument that, although generally significant in the first stages, did not qualitatively change any of our results.[117]

Taken as whole Table 4.7 strongly suggests that China is a good example of a low wage country trade shock. Import competition from low wage countries appears to stimulate faster technical change, whereas import competition from richer countries does not. According to our model, this is because imports from the South make the production of low-tech goods less

[117] For example in column (6) of Table 7 the coefficient (standard error) on trade weighted exchange rates was 0.391(0.178) in the first stage for IT and the coefficient on imports in the second stage remained insignificant (actually falling to -0.095 with a standard error of 0.172). For TFP the first stage coefficient (standard error) was 0.819(0.220) and the imports variable remained significant and positive in the second stage with a coefficient (standard error) of 0.210(0.081). For patents the first stage was very weak due to much fewer degrees of freedom with a coefficient (standard error) on the instrument of 0.082. The second stage coefficient on imports was negative but very imprecisely determined: -2.310(4.392).

profitable and increases incentives to move up the quality ladder. Rich country imports are more likely to be higher tech goods that for Schumpeterian reasons shrink profit margins and offset any pro-innovation effects of competition.

4.72 Heterogeneity: The effect of Chinese imports on innovation is stronger when there are more "trapped factors"

In Section II we suggested that firms with "trapped factors" (e.g. due to firm-specific human capital) may be less likely to innovate until a shock such as the reduction of trade barriers against Chinese goods lowers the opportunity cost of innovating. To test this idea we allow the effect of Chinese imports to be heterogeneous with respect to environments where we might think that trapped factors are more important.

Our first simple test is to construct a measure of the industry-specific wage premia that our model suggests is the product-specific human capital following an innovation (see Appendix A). We estimate these three digit inter-industry wage differentials in the standard way (e.g. Krueger and Summers, 1988) from a Mincerian wage equation using individual-level data. We do this for the UK as (i) there is abundant publicly available micro-data and (ii) we want to avoid conflating institutional constraints (like unions and minimum wages) with the underlying technology of the industry and the UK has the least regulated labor market of our European countries[118].

[118] For example, the OECD (2009) index of "strictness of employment protection in 2008" gives the UK the lowest score (i.e. highest flexibility) of 1.1 (on a scale of zero to 6) of all 30 developed countries with the exception of the US. By contrast, Portugal had the greatest degree of job protection with a score of 4.2.

In column (1) of Table 4.8 we repeat our basic patents equation. In column (2) we include our proxy for trapped factors, the measure of the industry-specific wage premium. This has a negative and significant correlation with innovation as the model would suggest as the opportunity cost of innovating is higher for firms with more trapped factors. Column (3) includes the key term: an interaction of the growth of Chinese imports and the industry wage premium. The coefficient on this term is positive and significant implying that the effect of Chinese competition is greater when there is more industry-specific human capital as the model predicts.

Using industry wage premium interprets the theory quite literally and it may be that trapped factors are a more general phenomenon. An alternative measure of trapped factors is to use measured TFP ("MFP") as a higher value of this term will reflect the fact that some firms have higher TFP than others (see Appendix A). The advantage of this measure is that it is firm specific, but a disadvantage is that we can only construct TFP for a sub-sample of the data. Column (4) of Table 4.8 presents the patent equations for this sub-sample. Even though the sample is smaller, the effect of Chinese import competition on patents is similar to that in the overall sample in Table 1 (0.284 vs. 0.321). We then include the firm's initial TFP in column (5) which, in line with the trapped factor model, is negatively correlated with subsequent patent growth. Column (6) includes the key interaction term between import growth and initial TFP. There is a significant and positive interaction suggesting that high TFP firms are more likely to respond by innovating when faced by a Chinese import shock than low productivity firms. This result has the same flavor as Aghion et al (2005) that the innovation in firms nearer the technology frontier responds more positively to competition, than low TFP firms. Unlike Aghion et al, however, we find no evidence of an inverted "U" which may be because we focus on competition from less developed countries who are near the bottom of the quality ladder, rather than an increase in general competition.[119]

We could not find any evidence that larger firms responded more to Chinese imports. But Holmes and Stevens (2010) argue that size is not an adequate proxy for productivity, finding that small plants actually do relatively better than larger plants following an increase in Chinese import competition. In their model, small firms survive by operating in product niches rather than the standardized products competing with China. Like Holmes and Stevens (2010) we find that size *per se* is an inadequate proxy for productivity, but document a new result that firms endogenously create niche products through innovation when faced by Chinese competition.

In summary, Table 4.8 as a whole suggests some heterogeneity of the effect of trade on innovation in a direction consistent with our simple trapped factor model.

4.73 Skill demand

To examine skill demand we use the UK Labor Force Survey, as none of our micro datasets has plant or firm level skills measures. We create a three-digit panel on the share of the college-educated workers in the total wage-bill. Since the impact of China is relatively common across Europe, we think the UK results should be broadly representative.

In column (1) of Table 4.9, we see that Chinese imports are associated with an increase in the wage-bill share of college-educated workers, suggesting Chinese trade raises the

---

[119] In a similar vein, Amiti and Khandelwal (2010) find stronger effects of trade on quality upgrading for firms closer to the quality frontier. Following Khandelwal (2010) we tried interacting imports with his average length of a quality ladder in the industry. The interactions typically went in the expected direction, but were insignificant.

demand for skills. In column (2) we see the standard result that IT is also associated with an increase in the share of wages for college workers. Including both variables into the regression in column (3) shows that both IT and Chinese imports are significant, although both have lower coefficients, suggesting part of the association of IT with skilled workers may be a proxy for the impact of developing country trade[120]. In column (4) we re-estimate this specification by OLS using the textile and apparel sample, and in column (5) report the IV results that support a causal impact of Chinese import competition on the demand for skilled workers. This is consistent with the model that Chinese trade leads firms to switch from producing older low-tech goods to the design and manufacture of new goods, which is likely to increase the demand for skilled workers.

4.74 Offshoring

We have focused on China's effect through competition in the final goods market, but an alternative way in which China could affect technical progress is through allowing Western firms to buy cheaper intermediate inputs and offshore low value added parts of the production chain.[121] We investigate this by adapting the offshoring measure of Feenstra and Hansen (1999) for China, which uses the input-output tables to measure for each industry the share of Chinese

[120] When disaggregating the wage bill share in relative wages and relative employment we find a positive association of Chinese imports with both components, but the strongest impact is on relative employment rather than relative wages.

[121] Intermediate inputs are stressed (in a developing country context) by Amiti and Konings (2006) and Goldberg et al, 2010b).

inputs in total imported inputs[122]. Column (1) of Table 4.10 includes this China offshoring measure in the patent equation. It enters with a positive but insignificant coefficient. Interestingly, in columns (2) and (3) we look at IT and TFP and *do* find a significant positive impact of offshoring. Throughout Table 4.10, the share of Chinese imports in the final goods market (our baseline measure) remains positive and significant with only slightly lower coefficients.[123]

We also investigated using the WTO quasi-experiment of Table 4.2 to construct "input quotas" using the input-output tables to calculate predicted falls in the barriers to using Chinese inputs. Looking at the reduced forms for the technology equations (i.e. simply regressing the five year growth of each technology measure on input quotas and country dummies interacted with time dummies), removal of input quotas had no significant impact on patents, but significantly increased IT intensity and TFP (see Table 4.10). When output quotas were also included in this specification, input quotas remained significant at the 5% level for the TFP equation, but were only significant at the 10% level for the IT equation.[124] Output quotas remained positive and significant in all three specifications.

---

[122] See Appendix B for details. We also considered the share of total imported inputs in all inputs (or all costs) like Feenstra-Hansen, but as with our analysis of total imports in the final goods market, it is the Chinese share (reflecting low wage country inputs) that is the dominant explanatory factor.

[123] This is compared to the baseline results in columns (1), (2) and (4) in Table 1 for patents, IT and TFP. The coefficient estimates in Table 10 imply a one standard deviation increase in offshoring has a similar marginal effect on IT and TFP (0.014 and 0.008 respectively) to a one standard deviation increase in Chinese imports (0.014 and 0.007 respectively).

[124] The coefficients (standard errors) on input quotas were 0.727(0.523), 0.696(0.365) and 0.290(0.136) in the patents, IT and TFP equations. We estimate these equations on industries where at least 0.5% of imported inputs are from China.

Together these results suggest that while offshoring does not increase overall innovation (as measured by patents) it does increase IT intensity and productivity, presumably since offshoring moves the less IT intensive and lower productivity parts of the production process overseas to China.

We re-estimated all the technology, employment and survival equations including extra terms in Chinese offshoring (see Table 4.A8). As with Table 4.10, these terms made little difference to the main patents equation but did have some effect on the IT and TFP equations, suggesting more of a role for offshoring in increasing the reallocation effects of China, broadly in line with the compositional models of sub-section IIC. We re-calculated the aggregate magnitudes of the effects of China on technical change including the offshoring terms (see Table 4.A9, the analog of Table 4.5). Although the overall effects on patents are not much changed (China still accounted for just under 15% of the increase in patenting), the implied effects of China on aggregate IT and TFP more than doubled suggesting that offshoring magnifies the product market competition effects of Chinese trade we have focused on. This implies that if anything, we are *underestimating* the effect of China by focusing on the final goods markets effects.

4.75 Product and industry switching

A leading compositional theory we discussed in the theory section was that in the face of Chinese import competition European firms change their product mix. To investigate this we examine whether a plant changes its primary four-digit industrial sector in the HH data, which has accurate four-digit industry data going back to 1999 (the other datasets have less reliable information on the changes in industry affiliation). On average 11% of plants switch industries over a five-year period, a substantial number that is consistent with evidence from recent papers.[125]

Table 4.11 begins by regressing a dummy for switching on Chinese imports and the usual controls, finding plants in industries exposed to China were more likely to switch industries. Column (2) includes a control for lagged IT intensity that reduces the probability of switching, but only slightly reduces the coefficient on Chinese imports. Column (3) includes employment growth, which has little impact. The second half of Table 4.11 uses IT intensity growth as the dependent variable. Column (4) shows that switching is indeed associated with greater use of IT, but the magnitude of the effect is small: plants who switched industries had a 2.5% faster growth in IT intensity than those who did not. Column (5) displays the standard regression for this sample, showing the positive relationship between IT intensity and Chinese imports for the sub-sample where we have switching data. Most importantly, column (6) includes the switching dummy; this reduces the coefficient on Chinese imports, but only by a small amount. A similar story is evident when we include employment growth in the final column. So industry switching is statistically significant but cannot account for much of Chinese import effects.

One limitation of this analysis is that our data does not allow us to observe product switching at a more disaggregate level. Bernard et al (2010, Table 4.5) show, however, that in US manufacturing firms three quarters of the firms who switched (five-digit) products did so across a four-digit industry. If we run column (5) on those plants who did not switch

[125] For example, Bernard, Redding and Schott (2010) on the US, Goldberg et al (2010a, b). Bernard et al (2006) found that 8% of their sample of US manufacturing plants switched four-digit industries over a five-year period.

industries, the Chinese imports effect remains strong (0.474 with a standard error of 0.082). This could still conceivably be driven by the small percentage of plants who switched five-digit sector within a four sector, but it seems unlikely given the small effect of controlling for four-digit switching on the Chinese imports coefficient.

4.76 Exports to China

We have focused on imports from China as driving changes in technology but as discussed in Section 6.2, exports may also have an impact through market size effects. Comtrade allows us to construct variable reflecting exports to China (as a proportion of total exports in the industry-country pair) in an analogous way to imports. Table 4.A10 presents the results, and shows that in every column of results exports are not significant. This is unsurprising as most of the theories of export-led productivity growth focus on exporting to *developed* countries rather than emerging economies, like China. It is unclear what benefit there is to learning, for example, from China that is usually thought of as being behind the European technology frontier. And in terms of market size, China’s share of the total world exports produced by European manufacturers is still relatively small at around 1.3%, so is not likely to drive technology change in the North.

4.8 Conclusions

In this paper we have examined the impact of trade on technical change in twelve European countries. Our motivation is that the rise of China which constitutes perhaps the most important exogenous trade shock from low wage countries to hit the “Northern” economies. This helps identify the trade-induced technical change hypothesis. We use novel firm and plant level panel data on innovation (patents and R&D), information technology, TFP and management practices combined with four-digit industry-level data on trade.

The results are easy to summarize. Our primary result is that the absolute volume of innovation as measured by patenting (and R&D) rose *within* firms who were more exposed to increases in Chinese imports. A similar large within firm effect is observe for other indicators of technical change such as TFP, IT intensity and management quality. Second, in sectors more exposed to Chinese imports, jobs and survival rate fell in low-tech firms (e.g. lower patenting intensity), but high-tech firms were are relatively sheltered (the between firm effect). Both within and between firm effects generate aggregate technological upgrading.

These results appear to be robust to many tests, including treating imports as endogenous using China's accession to the World Trade Organization in 2001. In terms of magnitudes, China could account for around 15% of the overall technical change in Europe between 2000 and 2007. This effect appears to be increasing over time and may even be an underestimate as we also identify a similar sized role for offshoring to China in increasing TFP and IT adoption (although not for innovation). This suggests that increased import competition with China has caused a significant technological upgrading in European firms in the affected industries through both faster diffusion and innovation. In terms of policy, our results imply that reducing import barriers against low wage countries like China may bring important welfare gains through technical change, subject to the caveats over general equilibrium effects discussed in sub-section 6.4.

There are several directions this work could be taken. First, we would like to investigate more deeply the impact of low wage countries on the labor market, using worker level data on the non-employment spells and subsequent wages of individuals most affected by Chinese trade. Much of the distributional impact depends on the speed at which the reallocation process takes place. Second, we want to complement our European analysis with a similar exercise in the US and other countries. Thirdly, we would like to further develop our trapped factor model, to see how important it is in explaining trade effects compared to the more conventional market size and competition effects. Finally, it would be helpful to more structurally extend the analysis to properly take into account general equilibrium effects. These areas are all being actively pursued.

REFERENCE LIST

Acemoglu, Daron (2002) “Technical Change, Inequality and the Labor Market”, *Journal of Economic Literature,* 40(1): 7-72.

Acemoglu, Daron (2008) “Equilibrium Bias of Technology”, *Econometrica,* 75(5): 1371-1409.

Acharya, Ram and Wolfgang Keller (2008) “Estimating the Productivity, Selection and Technology Spillover Effects of Imports”, NBER Working Paper No. 14079.

Aghion, Philippe and Gilles Saint-Paul (1998) “Virtues of Bad Times”, *Macroeconomic Dynamics*, 2: 322-344.

Aghion, Philippe, Nicholas Bloom, Richard Blundell, Rachel Griffith and Peter Howitt (2005) “Competition and Innovation: An Inverted U Relationship”, *Quarterly Journal of Economics,* 120(2): 701-728.

Amiti, Mary, and Caroline Freund (2010) “An Anatomy of China's Export Growth”, in Robert Feenstra and Shang-Jin Wei (eds) *China's Growing Role in World Trade,* Chicago Press.

Amiti, Mary, and Jozef Konings (2006) “Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia”, *American Economic Review,* 97(5): 1611-1638.

Amiti, Mary and Amit Khandelwal (2010) “Import Competition and Quality Upgrading”, University of Columbia mimeo.

Arkolakis, Costas, Peter Klenow, Svelenta Demidova and Andreas Rodriguez-Clare (2008) “Endogenous Variety and the Gains from Trade” *American Economic Review Papers and Proceedings,* 98 (4), 444-450

Arkolakis, Costas, Arnaud Costinot and Andres Rodríguez-Clare (2010) “New Trade Models, Same Old Gains”, MIT mimeo, forthcoming *American Economic Review*

Arrow, Kenneth (1962) “The Economic Implications of Learning by Doing”, *Review of Economic Studies*, 29(3):155-73.

Atkeson, Andrew and Ariel Burstein (2010) “Innovation, Firm Dynamics and International Trade”, *Journal of Political Economy*, 118(3), 433-484

Auer, Raphael and Andreas Fischer (2010) “The Effect of Low Wage Import Competition on US Inflationary Pressure”, *Journal of Monetary Economics,* 57(4),491-503,

Autor, David, Lawrence Katz and Alan Krueger (1998) “Computing Inequality: Have Computers Changed the Labor Market?” *Quarterly Journal of Economics,* 113(4): 1169-1213.

Autor, David, Dorn, David and Hansen, Gordon (2012) “The China Syndrome: local labor effects of import competition in the United States”, MIT Working Paper.

Bailey, Martin, Charles Hulten and David Campbell (1992) “Productivity Dynamics in Manufacturing Firms”, *Brookings Papers on Economic Activity: Microeconomics*, 4: 187-267.

Barlevy, Gadi (2007) “On the Cyclicality of Research and Development”, *American Economic Review*, 97(4): 1131-1164.

Bartel, Ann, Casey Ichinowski and Kathryn Shaw (2007) “How Does Information Technology Really Affect Productivity? Plant-level Comparisons of Product Innovation, Process Improvement and Worker Skills”, *Quarterly Journal of Economics,* 122 (4): 1721-1758.

Belenzon, Sharon and Tomer Berkovitz (2010) “Innovation in Business Groups”, *Management Science* 56(3): 519-535.

Bernard, Andrew and Bradford Jensen (2002) “The Death of Manufacturing Plants”, NBER Working Paper No. 9026.

Bernard, Andrew, Bradford Jensen and Peter Schott (2006) “Survival of the Best Fit: Exposure to Low-wage Countries and the (uneven) Growth of US Manufacturing Establishments”, *The Journal of International Economics*, 68(1): 219-237.

Bernard, Andrew, Stephen Redding and Peter Schott (2007) “Comparative Advantage and Heterogeneous Firms”, *Review of Economic Studies,* 74(1): 31-66.

Bernard, Andrew, Stephen Redding and Peter Schott (2010) “Multi-product Firms and Product Switching”, *American Economic Review,* 100(1): 70-97.

Bertrand, Marianne (2004) “From Invisible Handshake to the Invisible Hand? How Import Competition Changes the Employment Relationship”, *Journal of Labor Economics*, 22(4):723-765.

Bloom, Nicholas and John Van Reenen (2007) “Measuring and Explaining Management Practices across Firms and Countries”, *Quarterly Journal of Economics*, 122(4), 1341-1408.

Bloom, Nicholas and John Van Reenen (2010) “Why do Management Practices Differ across Firms and Countries?” *Journal of Economic Perspectives* 24(1) 203-224

Bloom, Nicholas, Mirko Draca and John Van Reenen (2011) “Trade Induced Technical Change? The impact of Chinese imports on innovation, IT and productivity” Centre for Economic Performance Working Paper No. 1000

Bloom, Nicholas, Rachel Griffith and John Van Reenen (2002) “Do R&D Tax Credits Work?” *Journal of Public Economics,* 85: 1-31.

Bloom, Nicholas, Sadun, Raffaella and Van Reenen, John (2009) “The Organization of Firms across Countries”, NBER Working Paper No 937

Bloom, Nicholas, Paul Romer, Stephen Terry and John Van Reenen (2012) “A Trapped Factor Model of Innovation”, LSE/Stanford mimeo http://www.stanford.edu/~nbloom/TF.pdf

Blundell, Richard, Rachel Griffith and John Van Reenen (1999) “Market Share, Market Value and Innovation”, *Review of Economic Studies* 66(3): 529-554.

Blundell, Richard, Griffith, Rachel and Windmeijer, Frank (2002) “Individual Effects and Dynamic Feedback in Count Data Models” *Journal of Econometrics*, 108, 113-131

Blundell, Richard, Amanda Gosling, Hidehiko Ichimura and Costas Meghir (2007) “Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds” *Econometrica*, 75 (2) 323-363

Brambilla, Irene, Amit Khandewal and Peter Schott (2010) “China’s Experience under the Multi-Fiber Agreement (MFA) and the Agreement on Textile and Clothing (ATC)”, in Robert Feenstra and Shang-Jin Wei (eds) *China's Growing Role in World Trade*, Chicago: Chicago University Press.

Broda, Christian and John Romalis (2009) “The Welfare Implications of Rising Price Dispersion”, University of Chicago mimeo.

Bugamelli, Matteo, Fabiano Schivardi and Roberta Zizza (2008) “The Euro and Firm Restructuring”, NBER Working Paper No. 14454.

Bustos, Paula (2011) “Trade Liberalization, Exports and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinean Firms”, *American Economic Review*, 101(1), 304-40,

Cameron,Colin and Trivedi, Pravin (2005) *Microeconometrics: Methods and Applications. New York: Cambridge University Press.*

Card, David (2001) “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration”, *Journal of Labor Economics*, 19(1): 22-64.

Costantini, James and Marc Melitz (2008) “The Dynamics of Firm-Level Adjustment to Trade Liberalization” in Elhanan Helpman, Dalia Marin and Thierry Verdier *The organization of firms in a global Economy*, Cambridge: Harvard University Press

Coe, Daniel and Elhanan Helpman (1995) “International R&D Spillovers”, *European Economic Review*, 39(5): 859-887.

Davis, Steven and Haltiwanger, John (1992) “Gross Job Creation, Gross Job Destruction, and Employment Reallocation”. *The Quarterly Journal of Economics*, 107(3), 819-63.

De Loecker, Jan (201) “Product Differentiation, Multi-product Firms and Estimating the Impact of Trade Liberalization on Productivity”, *Econometrica* 79(5) 1407–1451

DiNardo, John, Nicole Fortin, and Thomas Lemieux (1996) "Labor Market Institutions and the Distribution of Wages, 1973-1992". *Econometrica,* LXV: 1001-1044.

Dunne, Timothy, Shawn Klimek and James Schmitz (2008) "Does Foreign Competition Spur Productivity? Evidence from Post WWII US Cement Manufacturing", Minneapolis Fed mimeo.

Eaton, Jonathan and Samuel Kortum (1999) "International Technology Diffusion: Theory and Measurement" *International Economic Review*, 40(3), 537-70

Eaton, Jonathan and Samuel Kortum (2001), "Technology, Trade, and Growth: A Unified Framework" *European Economic Review*, 45(4-6), 742-755.

Eaton, Jonathan and Samuel Kortum (2002) "Technology, Geography, and Trade", *Econometrica*, 70(5), 1741-1779

Eslava, Marcela, John Haltiwanger and Adriana Kugler (2009) "Trade Reforms and Market Selection: Evidence from Manufacturing Plants in Columbia", NBER Working Paper 14935

Feenstra, Robert, Robert Lipsey, Haiyan Deng, Alyson Ma and Hengyong Mo (2005) "World Trade Flows, 1962-2000", NBER Working Paper, No. 11040.

Feenstra, Robert and Gordon Hansen (1999) "The Impact of Outsourcing and High-technology Capital on Wages: Estimates for the US, 1979-1990", *Quarterly Journal of Economics,* 114(3): 907-940.

Freeman, Richard and Morris Kleiner (2005) "The Last American Shoe Manufacturers", *Industrial Relations*, 44: 307-342.

Goldberg, Pinelopi, Amit Khandewal, Nina Pavcnik and Petia Topalova (2010a) "Multi-product Firms and Product Turnover in the Developing World: Evidence from India." *Review of Economics and Statistics*, 92(4), 1042-1049

Goldberg, Pinelopi, Amit Khandewal, Nina Pavcnik and Petia Topalova (2010b) "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India", *Quarterly Journal of Economics*, 125(4), 1727-1767

Grossman, Gene and Elhanan Helpman (1991) "Quality Ladders and Product Cycles", *Quarterly Journal of Economics*, 106: 557-586.

Grossman, Gene and Elhanan Helpman (1992) *Innovation and Growth in the Global Economy*, Cambridge: MIT Press.

Grossman, Gene and Estaban Rossi-Hansberg (2008) "Trading Tasks: A Simple Theory of Offshoring", *American Economic Review*, 98(5): 1978-1997.

Hausman, Jerry; Hall, Bronwyn; and Griliches, Zvi (1984) "Econometric Models of Count Data with an Application to the Patents-R&D Relationship", *Econometrica*, 52(4):909-938.

Holmes, Thomas and John Stevens (2010) "An Alternative Theory of the Plant Size Distribution with an Application to Trade", mimeo University of Minnesota.

Holmes, Thomas, David Levine and James Schmitz (2008) "Monopoly and the Incentives to Innovate When Adoption Involves Switchover Disruptions", Minneapolis Fed working paper.

Honoré Bo and Kyriazidou, Ekaterini (2000) "Panel Data Discrete Choice Models with Lagged Dependent Variables" *Econometrica*, 68, 839-874

Iacovone, Leonardo and Beata Javorcik (2008) "Multi-product exporters: Diversification and Micro-level Dynamics" World Bank Policy Research Working Paper No. 4723

Jones, Charles and John Williams. (1998) "Measuring the Social Rate of Return to R&D", *Quarterly Journal of Economics*, 113(4), pp. 119-35

Khandelwal, Amit (2010) "The Long and Short (of) Quality Ladders", *Review of Economic Studies*, 77(4), 1450-1476

Klette, Tor Jakob and Zvi Griliches (1996) "The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous", *Journal of Applied Econometrics*, 11(4): 343-361

Krueger, Alan and Lawrence Summers (1988) "Efficiency Wages and the Inter-industry Wage Structure" *Econometrica*, 56(2), 259-93

Krugman, Paul (1980) "Scale Economies, Product Differentiation, and the Pattern of Trade", *American Economic Review*, 70(5): 950-959.

Krugman, Paul (2008) "Trade and Wages, Reconsidered", *Brookings Papers on Economic Activity*, Spring, 103-154

Kyriazidou, Ekaterini (1997) "Estimation of a Panel Data Sample Selection Model", *Econometrica*, 65(6), 1335-1364.

Levchenko, Andrei and Jing Zhang (2010) "The Evolution of Comparative Advantage: Measurement and Welfare Implications", University of Michigan mimeo

Lelarge, Claire and Benjamin Nefussi (2008) "Exposure to Low-wage Competition, Activity Changes and Quality Upgrading: An Empirical Assessment'', OECD mimeo.

Lileeva, Alla and Dan Trefler (2010) "Improved Access to Foreign Markets Raises Plant-Level Productivity ... for Some Plants", *Quarterly Journal of Economics*, 125(3).

Machin, Stephen and John Van Reenen (1998) "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries", *Quarterly Journal of Economics*, 113(4): 1215-1244.

Manski, Charles (1994): "The Selection Problem," in *Advances in Econometrics, Sixth World Congress*

Manski, Charles and J. Pepper (2000): "Monotone Instrumental Variables: With Application to the Returns to Schooling," *Econometrica*, 68, 997–1010.

Melitz, Marc (2003) "The Impact of Trade on Intra-industry Reallocations and Aggregate Productivity Growth", *Econometrica*, 71: 1695-1725.

OECD (2009) *OECD Economic Surveys: Ireland, Figure 3.3.* OECD, Paris, France

Olley, Steven and Ariel Pakes (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry", *Econometrica* 64: 1263–1297.

Ossa, Ralph and Chang-Tai Hsieh (2010) "A Global View of Productivity Growth in China", mimeo, University of Chicago

Pavcnik, Nina (2002) "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Establishments", *The Review of Economic Studies*, 69(1): 245-276.

Pierce, Justin and Peter Schott (2010) "Concording US Harmonized System Codes Over Time", Mimeo, Yale University.

Raith, M. (2003) "Competition, Risk and Managerial Incentives" *American Economic Review*, 93, 1424-1436

Rivera-Batiz, Luis and Paul Romer (1991) "Economic Integration and Endogenous Growth" *Quarterly Journal of Economics*, 106(2), 531-55

Schmidt, Klaus (1997) "Managerial Incentives and Product Market Competition" *Review of Economic Studies* 64(2) 191-213

Schmookler, Jacob (1966) *Invention and Economic Growth*, Cambridge: Harvard University Press.

Teshima, Kensuke (2008) "Import Competition and Innovation at the Plant Level: Evidence from Mexico", Columbia University mimeo.

Thoenig, Matthias and Thierry Verdier (2003) "A Theory of Defensive Skill Biased Innovation and Globalization", *American Economic Review*, 93(3): 709-728.

Trefler, Daniel (2004) "The Long and Short of the Canada-US Free Trade Agreement", *American Economic Review*, 94: 870-895.

Verhoogen, Eric (2008) "Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector", *Quarterly Journal of Economics*, 123(2): 489-530.

Wooldridge, Jeff (1995) "Selection Corrections for Panel Data Models under Conditional Mean Independence Assumptions*," Journal of Econometrics* 68, 115-132

Wood, Adrian (1994) *North-South Trade, Employment and Inequality:* Changing *Fortunes in a Skill Driven World,* Oxford: Clarendon Press.

Yeaple, Stephen (2005) "Firm Heterogeneity, International Trade and Wages", *Journal of International Economics*, 65(1): 1-20.

4.9 Tables for Chapter 4.

TABLE 4.1: TECHNICAL CHANGE WITHIN INCUMBENT FIRMS AND PLANTS

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	**Δln(PATENTS)**	**Δln(IT/N)**	**Δln(R&D)**	**ΔTFP**	**ΔMANAGEMENT**
Estimation method	**5 year diffs**	**5 year diffs**	**5 year diffs**	**5 year diffs**	**3 year diffs**
Change in Chinese Imports	0.321***	0.361**	1.213**	0.257***	0.814***
ΔIMP_{jk}^{CH}	(0.102)	(0.076)	(0.549)	(0.072)	(0.314)
Sample period	2005-1996	2007-2000	2007-1996	2005-1996	2010-2002
Number of Units	8,480	22,957	459	89,369	1,576
Number of country by industry clusters	1,578	2,816	196	1,210	579
Observations	30,277	37,500	1,626	292,167	3,607

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country by four-digit industry pair in parentheses (except columns (3) and (5) which are three-digit industry by country). All changes are in five-year differences, e.g. ΔIMP_{jk}^{CH} represents the 5-year difference in Chinese imports as a fraction of total imports in a four-digit industry by country pair (except column (5) which is in three-year long differences). All columns include a full set of country by year dummies. *Δln(PATENTS)* is the change in *ln(1+PAT), PAT* = count of patents. *IT/N* is the number of computers per worker. *R&D* is expenditure on research and development. TFP is estimated using the de Loecker (2011) version of the Olley-Pakes (1996) method separately for each industry based on 1.4m underlying observations (see Appendix C) and *Management* is the average score on the 18 Bloom and Van Reenen (2007) management questions around monitoring, targets and incentives. The 12 countries include Austria, Denmark, Finland, France, Germany, Ireland, Italy, Norway, Spain, Sweden, Switzerland and the UK for all columns except (4) which only includes France, Italy, Spain and Sweden (the countries where we have good data on intermediate inputs) and column (5) which covers France, Germany, Italy, Ireland, Sweden and the UK. Dummies for establishment type (Divisional HQ, Divisional Branch, Enterprise HQ or a Standalone Branch) are included in column (2). Standard survey noise controls such as interviewer dummies and interview/interviewee controls (e.g. tenure in firm) are included in column (5) as in Bloom and Van Reenen (2007). Units are firms in all columns except (2) where it refers to plants.

TABLE 4.2: CONTROLING FOR UNOBSERVED TECHNOLOGY SHOCKS

PANEL A: USING CHANGES IN QUOTAS AS AN IV (TEXTILE AND APPAREL INDUSTRIES ONLY)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	PATENTING ACTIVITY			INFORMATION TECHNOLOGY			TOTAL FACTOR PRODUCTIVITY		
Dependent Variable:	**Δln(PATENTS)**	**ΔIMP^{CH}**	**Δln(PATENTS)**	**Δln(IT/N)**	**ΔIMP^{CH}**	**Δln(IT/N)**	**ΔTFP**	**ΔIMP^{CH}**	**ΔTFP**
Method:	**OLS**	**First Stage**	**IV**	**OLS**	**First Stage**	**IV**	**OLS**	**First Stage**	**IV**
Change Chinese Imports	1.160***		1.864*	1.284***		1.851***	0.620***		1.897**
	(0.377)		(1.001)	(0.172)		(0.400)	(0.100)		(0.806)
Quotas removal		0.108***			0.088***			0.068***	
		(0.022)			(0.019)			(0.026)	
Sample period	2005-1999	2005-1999	2005-1999	2005-2000	2005-2000	2005-2000	2005-1999	2005-1999	2005-1999
Number of units	1,866	1,866	1,866	2,891	2,891	2,891	55,791	55,791	55,791
Number industry clusters	149	149	149	83	83	83	187	187	187
Observations	3,443	3,443	3,443	2,891	2,891	2,891	55,791	55,791	55,791

PANEL B: USING "INITIAL CONDITIONS" AS AN INSTRUMENTAL VARIABLE (ALL INDUSTRIES)

Dependent Variable	**Δln(PATENTS)**	**ΔIMP^{CH}**	**Δln(PATENTS)**	**Δln(IT/N)**	**ΔIMP^{CH}**	**Δln(IT/N)**	**ΔTFP**	**ΔIMP^{CH}**	**ΔTFP**
Method:	**OLS**	**First Stage**	**IV**	**OLS**	**First Stage**	**IV**	**OLS**	**First Stage**	**IV**
Change in Chinese Imports	0.321***		0.495**	0.361***		0.593***	0.257***		0.507*
	(0.117)		(0.224)	(0.106)		(0.252)	(0.087)		(0.283)
Chinese imports in SIC4*US		0.167***			0.124***			0.078***	
&EU Chinese import growth		(0.017)			(0.002)			(0.021)	
Sample period	2005-1996	2005-1996	2005-1996	2007-2000	2007-2000	2007-2000	2005-1996	2005-1996	2005-1996
Number of Units	8,480	8,480	8,480	22,957	22,957	22,957	89,369	89,369	89,369
Number of industry clusters	304	304	304	371	371	371	354	354	354
Observations	30,277	30,277	30,277	37,500	37,500	37,500	292,167	292,167	292,167

(CONTINUED OVER)

(TABLE 4.2, CONTINUED)

PANEL C: INCLUDE INDUSTRY TRENDS (OLS, ALL INDUSTRIES)

Dependent Variable	(1) **Δln(PATENTS)**	(2) **Δln(PATENTS)**	(3) **Δln(IT/N)**	(4) **Δln(IT/N)**	(5) **ΔTFP**	(6) **ΔTFP**
Change in Chinese Imports	0.321***	0.191*	0.361***	0.170**	0.257***	0.128**
	(0.102)	(0.102)	(0.076)	(0.082)	(0.072)	(0.053)
Three Digit Industry trends?	No	Yes	No	Yes	No	Yes
Sample period	2005-1996	2005-1996	2007-2000	2007-2000	2005-1996	2005-1996
Number of Units	8,480	8,480	22,957	22,957	89,369	89,369
Number of clusters	1,578	1,578	2,816	2,816	1,210	1,210
Observations	30,277	30,277	37,500	37,500	292,167	292,167

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. In all panels we use the same specifications as Table 1 columns (1), (2) and (4) but estimate by instrumental variables (IV). In Panel A the IV is "Quota removal" is based on EU SIGL data and defined as the (value weighted) proportion of HS6 products in the four-digit industry that were covered by a quota restriction on China in 2000 (prior to China's WTO accession) that were planned to be removed by 2005 (see the Appendix D for details). Sample only includes textiles and apparel industries. In **Panel B** the IV is the share of Chinese imports in all imports in an industry across the whole of the Europe and the US (6 years earlier) interacted with the aggregate growth in Chinese imports in Europe. The base year is (t-6). **Panel C** reproduces the baseline OLS regressions in columns (1), (3) and (5) and then includes a full set of three-digit dummies in columns (2), (4) and (6). Since these specifications are in long differences this is equivalent to including three digit trends in the levels specification. The number of units is the number of firms in all columns except the IT specification where it is the number of plants. All columns include country by year effects. Standard errors for all regressions are clustered by four-digit industry in parentheses in panels A and B and by four-digit industry by country pairs in Panel C.

TABLE 4.3: EMPLOYMENT AND EXIT

PANEL A: EMPLOYMENT

Dependent Variable: Employment Growth, $\Delta \ln N$	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Technology variable (TECH)		Patent stock	Patent stock	Patent stock	IT	TFP
Sample	All Firms	All Firms	Patenting firms	Patenting firms	IT sample	TFP sample
Change in Chinese Imports	-0.349***	-0.352***	-0.361***	-0.434***	-0.379***	-0.382***
ΔIMP_{jk}^{CH}	(0.067)	(0.067)	(0.134)	(0.136)	(0.105)	(0.093)
Change in Chinese imports*technology at t-5		1.546**		1.434**	0.385**	0.956**
ΔIMP_{jk}^{CH} $*TECH_{t-5}$		(0.757)		(0.649)	(0.157)	(0.424)
Technology at t-5	0.513***	0.469***	0.389***	0.348***	0.230***	0.256***
$TECH_{t-5}$	(0.050)	(0.058)	(0.043)	(0.049)	(0.010)	(0.016)
Number of Units	189,563	189,563	6,335	6,335	22,957	89,369
Number of country by industry clusters	3,123	3,123	1,375	1,375	2,816	1,210
Observations	581,474	581,474	19,844	19,844	37,500	292,167

PANEL B: EXIT

Dependent Variable: *SURVIVAL*	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Technology variable		Patent stock	Patent stock	Patent stock	IT	TFP
Sample	All Firms	All Firms	Patenting firms	Patenting firms	IT sample	TFP sample
Change in Chinese Imports	-0.122***	-0.122***	-0.065	-0.089	-0.182**	-0.189***
ΔIMP_{jk}^{CH}	(0.036)	(0.036)	(0.047)	(0.050)	(0.072)	(0.056)
Change in Chinese imports*technology at t-5		0.391**		0.261**	0.137	0.097
ΔIMP_{jk}^{CH} $*TECH_{t-5}$		(0.018)		(0.114)	(0.112)	(0.076)
Technology at t-5	0.052***	0.040***	-0.006	-0.014	-0.002	-0.003
$TECH_{t-5}$	(0.008)	(0.011)	(0.007)	(0.009)	(0.006)	(0.004)
Survival Rate for Sample (mean)	0.929	0.929	0.977	0.977	0.886	0.931
Number of country by industry clusters	3,369	3,369	1,647	1,647	2,863	1,294
Observations (and number of units)	490,095	490,095	7,985	7,985	28,624	268,335

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation by OLS with standard errors (clustered by country by four-digit industry pair) in parentheses. ΔIMP^{CH} represents the 5-year difference in Chinese imports as a fraction of total imports in a four-digit industry by country pair. In columns (1) to (4) *TECH* is ln[(1+ the firm's patent stock)/employment]. In column (5) *TECH* is computers per employee *(IT/N)* and in column (6) *TECH* is *TFP*. 12 Countries in all columns except column (6) which is for four countries. In columns (3) and (4) only "patenting firms" (defined as a firm that had at least one European patent between 1978 and 2007) included. Sample period is 2005-1996 for all except column (5) which is 2007-2000. Number of units is the number of firms in all columns except (5) where it is the number of plants. All columns include country by year effects. **In Panel A** the dependent variable is the five year difference of ln(employment). **In Panel B** the dependent variable (*SURVIVAL*) refers to whether an establishment that was alive in 2000 was still alive in 2005 for the HH sample in column (5). In the other columns it is based on Amadeus company status (Appendix B) and is defined on the basis of whether a firm alive in 2000 was dead by 2005.

TABLE 4.4: APPROXIMATE MAGNITUDES

PANEL A: Increase in Patents per employee attributable to Chinese imports (as a % of the total increase over the period)

Period	Within	Between	Exit	Total
2000-07	5.8	6.3	2.5	14.7

PANEL B: Increase in IT per employee attributable to Chinese imports (as a % of the total increase over the period)

Period	Within	Between	Exit	Total
2000-07	9.8	3.1	1.2	14.1

PANEL C: Increase in Total Factor Productivity attributable to Chinese imports (as a % of the total increase over the period)

Period	Within	Between	Exit	Total
2000-07	8.1	3.4	0.3	11.8

Notes: Panel A reports the share of aggregate IT intensity accounted for by China, Panel B the increase in patents; and the Panel C the increase in total factor productivity. This is calculated by multiplying the regression coefficients and the observed Chinese import share growth to generate a predicted change in IT/Employee, Patents/Employee and TFP 2000 to 2007 inclusive. This aggregate predicted growth in IT/Employee is then divided by the average annual change in IT/employee between 1999 to 2007 (2.5%). The aggregate predicted change in Patents/Employee is then divided by 3.5% (the aggregate annual growth rate of patents from 1986 to 2006 in the USPTO) and the aggregate predicted growth in TFP is divided by 2% (the average TFP growth in manufacturing).

TABLE 4.5: COMPARING INDUSTRY LEVEL REGRESSIONS TO FIRM LEVEL REGRESSIONS

PANEL A. INDUSTRY-COUNTRY LEVEL

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent Variable:	**Δln(Prices)**	**Δln(Employment)**	**Δln(Profits /Sales)**	**Δln(PATENTS)**	**Δln(PATENTS)**	**Δln(IT/N)**	**Δln(IT/N)**	**Δln(R&D)**	**Δln(R&D)**	**Δln(TFP)**
Change in Chinese Imports ΔIMP_{jk}^{CH}	-0.447** (0.216)	-0.422*** (0.148)	-0.112** (0.052)	0.368 * (0.200)	0.368* (0.200)	0.399*** (0.120)	0.354*** (0.120)	2.145* (1.186)	1.791** (0.829)	0.326*** (0.072)
Change in employment					0.005 (0.012)		-0.088*** (0.013)			
Change in ln(Production)									-0.297 (0.403)	
Sample period	2006-2000	2005-1996	2007-2000	2005-1996	2005-1996	2007-2000	2007-2000	2007-2000	2007-2000	2005-1996
Country by industry clusters	131	2,990	2,295	1,646	1,646	2,902	2,902	151	151	1,140
Observations	262	11,800	5,372	6,888	6,888	7,409	7,409	322	322	5,660

PANEL B. FIRM LEVEL EQUIVALENT (ALLOCATING FIRM TO A SINGLE FOUR-DIGIT INDUSTRY)

Dependent Variable:	**Δln(Prices)**	**Δln(Employment)**	**Δln(Profits /Sales)**	**Δln(PATENTS)**	**Δln(PATENTS)**	**Δln(IT/N)**	**Δln(IT/N)**	**Δln(R&D)**	**Δln(R&D)**	**Δln(TFP)**
Change in Chinese Imports ΔIMP_{jk}^{CH}	*No firm-level price data available*	-0.280**** (0.066)	-0.043*** (0.008)	0.171** (0.082)	0.215** (0.098)	0.361** (0.076)	0.195*** (0.067)	1.213** (0.549)	1.545*** (0.330)	0.164*** (0.051)
Change in employment					0.015* (0.009)		-0.617*** (0.010)			
Change in ln(Production)									0.558*** (0.043)	
Years		2005-1996	2007-2000	2005-1996	2005-1996	2007-2000	2007-2000	2007-2000	2007-2000	2005-1996
Country by industry clusters		2,814	2,259	1,578	1,464	2,816	2,816	196	196	1,018
Observations		556,448	214,342	30,277	22,938	37,500	37,500	1,626	1,626	241,810

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Panel A uses data aggregated to the industry by country level and panel B is the firm level equivalent specification with firms allocated to a single industry (except columns (6) and (7) which are plant level). Coefficients estimated by OLS in five-year differences with standard errors (clustered by industry-country pair) in parentheses below coefficients. Chinese imports are measured by the value share of Chinese imports in total imports. There are 12 countries in all columns except (10) which only includes France, Italy, Spain and Sweden (where we have good data on intermediate inputs) and (3) which is based on Germany, France, Finland, France, Spain and Sweden (where gross profit information is available). All columns include country-year effects. The dependent variable in column (1) is producer prices and is measured at the two-digit level. In column (3) the dependent variable is (pre-tax and interest) profits rates. Columns (8) and (9) in Panel A use industry R&D data from the OECD STAN database and includes Germany, Denmark, Spain, Finland, France, the UK, Italy, Norway and Sweden, and is run at the two-digit level. In column (10) productivity is estimated using the de Loecker (2011) version of the Olley-Pakes method separately for each two-digit industry (see text). All firms are allocated to a single four-digit industry unless otherwise stated (i.e. we do not use the multiple-industry information exploited in the other tables) in order to make the two Panels comparable.

TABLE 4.6: ASSESSING DYNAMIC SELECTION BIAS IN THE PATENTS EQUATION

	(1)	(2)	(3)	(4)
Estimator	**5 year long differences**	**5 year long differences**	**Fixed effects Negative Binomial**	**Fixed effects Negative Binomial**
Method	**Baseline**	**Worst case Lower Bound**	**Baseline**	**Worst case Lower Bound**
Change in Chinese Imports $\Delta\left(M_{jk}^{China} / M_{jk}^{World}\right)$	0.321*** (0.102)	0.271*** (0.098)		
Level of Chinese Imports $\left(M_{jk}^{China} / M_{jk}^{World}\right)$			0.397*** (0.168)	0.389*** (0.165)
Number of Clusters	1,578	1,662	1,578	1,662
Number of Firms	8,480	8,732	8,480	8,732
Number of Observations	30,277	31,272	74,038	75,463

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Sample period is 1996-2005 for all columns. Estimation in columns (1) and (2) by OLS in long-differences and by Negative Binomial count data model with fixed effects using the Blundell et al (1999) technique in columns (3) and (4). Standard errors (clustered by country by four-digit industry pair) in parentheses. "Worst case lower bounds" impute a value of zero to all observations through 2005 where a firm dies (death is defined as in Table 3B). There are more observations for the Negative Binomial than five year long differences as we are using observations with less than five continuous years. All columns include a full set of country by year dummies. 12 countries included in all samples.

TABLE 4.7: LOW WAGE COUNTRY AND HIGH WAGE COUNTRY IMPORTS

PANEL A: DEPENDENT VARIABLE IS CHANGE IN LN(PATENTS)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Change in Chinese Imports	0.182**	0.063			0.182**		0.178**
$\Delta\left(M_{jk}^{China} / D_{jk}\right)$	(0.074)	(0.125)			(0.073)		(0.077)
Change in Non-China Low Wage Imports		0.152					
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$		(0.128)					
Change in All Low Wage Imports			0.106***				
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$			(0.040)				
Change in High Wage Imports				0.004	0.003		
$\Delta\left(M_{jk}^{High} / D_{jk}\right)$				(0.019)	(0.019)		
Change in World Imports						0.017	0.004
$\Delta\left(M_{jk} / D_{jk}\right)$						(0.018)	(0.018)
Number of Firms	8,364	8,364	8,364	8,364	8,364	8,364	8,364
Number of industry-country clusters	1,527	1,527	1,527	1,527	1,527	1,527	1,527
Number of Observations	29,062	29,062	29,062	29,062	29,062	29,062	29,062

PANEL B: DEPENDENT VARIABLE IS CHANGE IN IT INTENSITY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Change in Chinese Imports	0.129***	0.126***			0.128***		0.120***
$\Delta\left(M_{jk}^{China} / D_{jk}\right)$	(0.028)	(0.029)			(0.028)		(0.029)
Change in Non-China Low Wage Imports		0.018					
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$		(0.051)					
Change in All Low Wage Imports			0.127***				
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$			(0.025)				
Change in High Wage Imports				0.014	0.002		
$\Delta\left(M_{jk}^{High} / D_{jk}\right)$				(0.009)	(0.009)		
Change in World Imports						0.024***	0.007
$\Delta\left(M_{jk} / D_{jk}\right)$						(0.009)	(0.009)
Number of Units	20,106	20,106	20,106	20,106	20,106	20,106	20,106
Number of industry-country clusters	2,480	2,480	2,480	2,480	2,480	2,480	2,480
Number of Observations	31,820	31,820	31,820	31,820	31,820	31,820	31,820

PANEL C: DEPENDENT VARIABLE IS CHANGE IN TOTAL FACTOR PRODUCTIVITY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Change in Chinese Imports	0.065***	0.092**			0.071***		0.062**
$\Delta\left(M_{jk}^{China} / D_{jk}\right)$	(0.020)	(0.048)			(0.021)		(0.022)
Change in Non-China Low Wage Imports		-0.026					
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$		(0.041)					
Change in All Low Wage Imports			0.050***				
$\Delta\left(M_{jk}^{Low} / D_{jk}\right)$			(0.014)				
Change in High Wage Imports				0.007	-0.006		
$\Delta\left(M_{jk}^{High} / D_{jk}\right)$				(0.006)	(0.007)		
Change in World Imports						0.014**	0.002
$\Delta\left(M_{jk} / D_{jk}\right)$						(0.006)	(0.007)
Number of Firms	89,369	89,369	89,369	89,369	89,369	89,369	89,369
Number of industry-country clusters	1,210	1,210	1,210	1,210	1,210	1,210	1,210
Number of Observations	293,167	293,167	293,167	293,167	293,167	293,167	293,167

Notes: *** denotes 1% significance, ** denotes 5% significance, * denotes 10% significance. Estimation is by OLS with standard errors clustered by four-digit industry pair. $\Delta\left(M_{jk}^{China} / D_{jk}\right)$ represents the 5-year difference in Chinese imports normalized by domestic production (D). is the 5-year difference in All Low Wage Country imports normalized by domestic production (D) $\Delta\left(M_{jk}^{High} / D_{jk}\right)$ is the 5-year difference in total World Imports normalized by domestic production (D). Production data is from Eurostat's Prodcom database (no Swiss data). All specifications include country-year dummies. In Panel B we include site-type dummies and employment growth as additional controls. Sample period is 2000-2007 for Panel B and 1996-2005 for Panels A and C. 12 countries.

TABLE 4.8: HETEROGENEITY - THE CHINA EFFECT ON INNOVATION IS GREATER FOR FIRMS WITH MORE "TRAPPED FACTORS"

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable: Δln(PATENTS)						
Change Chinese Imports ΔIMP_{jk}^{CH}	0.321*** (0.102)	0.192*** (0.090)	0.202*** (0.092)	0.284* (0.157)	0.343** (0.153)	-2.466*** (0.848)
Industry wage premia		-0.343*** (0.065)	-0.411*** (0.069)			
Change Chinese Imports * Industry Wage premia			2.467*** (1.171)			
Total Factor Productivity TFP_{t-5}					-0.232*** (0.046)	-0.287*** (0.050)
Change Chinese Imports *TFP_{t-5} ΔIMP_{jk}^{CH} *TFP_{t-5}						1.464*** (0.462)
Number of units	8,480	8,480	8,480	5,014	5,014	5,014
Number of clusters	1,578	1,578	1,578	1,148	1,148	1,148
Number of Observations	30,277	30,277	30,277	14,500	14,500	14,500

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country and four-digit industry cell in parentheses. 12 countries. Industry wage premia defined as coefficients on three digit industry dummies in a wage regression implemented using the UK LFS pooled cross-sections from 1996-2008 (see Appendix A). The ln(hourly wage) regression includes controls for a quadratic in experience, schooling, region and gender. TFP is calculated in the same way as rest of paper using the de Loecker (201) method (see Appendix C)

TABLE 4.9: RELATIVE DEMAND FOR COLLEGE EDUCATED WORKERS INCREASES WITH CHINESE IMPORTS

	(1)	(2)	(3)	(4)	(5)
Dependent Variable:	**Δ(Wage bill Share of college educated)**	**Δ(Wage bill Share of college educated)**	**Δ(Wage bill Share of college educated)**	**Δ(Wage bill Share of college educated)**	**Δ(Wage bill Share of college educated)**
Sample	**All**	**All**	**All**	**Textiles & Clothing**	**Textile & Clothing**
Method	**OLS**	**OLS**	**OLS**	**OLS**	**IV**
Change in Chinese Imports, ΔIMP_{jk}^{CH}	0.144*** (0.035)		0.099** (0.043)	0.166*** (0.030)	0.227*** (0.053)
Change in IT intensity $\Delta \ln(IT/N)$		0.081** (0.024)	0.050* (0.026)		
F-test of excluded IV					9.21
Industry Clusters	72	72	74	17	17
Observations	204	204	204	48	48

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. The sample period is 1999-2006. The dependent variable is the five-year difference in the wage bill share of college-educated workers. Estimation is by OLS with standard errors clustered by three-digit industry pair in parentheses. This data is a three-digit industry panel for the UK between 2000 and 2007 (based on aggregating up different years of the UK Labor Force Survey). All manufacturing industries in columns (1) - (3) and textiles and clothing industries sub-sample in columns (4)-(5). IV regressions use Quota removal (the height of the quota in the three-digit industry in 2000 prior to China joining the WTO). All regressions weighted by number of observations in the Labor Force Survey in the industry cell. All regressions control for year dummies.

TABLE 4.10: OFFSHORING

Dependent Variable:	(1) Δln(PATENTS)	(2) Δln(IT/N)	(3) Δln(TFP)
Change in Chinese Imports	0.313***	0.279***	0.189***
ΔIMP_{jk}^{CH}	(0.100)	(0.080)	(0.082)
Change in Chinese Imports in source industries	0.173	1.685***	1.396***
$\Delta OFFSHORE$	(0.822)	(0.517)	(0.504)
Number of units	8,480	22,957	89,369
Number of industry-country clusters	1,578	2,816	1,210
Observations	30,277	37,500	292,167

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country and four-digit industry cell in parentheses. 12 countries except column (3) where there are four countries. "Number of units" represents the number of firms in all columns except (2) where it is plants. Offshoring is defined as in Feenstra and Hansen (1999) except it is for Chinese imports only, not all low wage country imports (see Appendix B).

TABLE 4.11: INDUSTRY/PRODUCT SWITCHING AND TECHNICAL CHANGE

Dependent Variable:	(1) SWITCHED INDUSTRY	(2) SWITCHED INDUSTRY	(3) SWITCHED INDUSTRY	(4) Δln(IT/N)	(5) Δln(IT/N)	(6) Δln(IT/N)
Change in Chinese imports	0.138***	0.132***	0.131***		0.469***	0.466***
ΔIMP_{jk}^{CH}	(0.050)	(0.050)	(0.050)		(0.083)	(0.083)
IT intensity (t-5)		-0.018**	-0.018**			
$(IT/N)_{t-5}$		(0.007)	(0.008)			
Industry Switching				0.025***		0.023*
				(0.012)		(0.012)
Employment growth			-0.002			
$\Delta \ln(Employment)$			(0.006)			
Observations	32,917	32,917	32,917	32,917	32,917	32,917

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. The plant-level Harte-Hanks data is used for all regressions reported in the table."Switched Industry" is a dummy variable equal to unity if a plant switched four-digit industry classification over the 5-year period. Estimation is by OLS standard errors clustered by four-digit industry and country. 12 Countries. All regressions include country-year effects and site-type controls. Sample period is 2000 to 2007.

4.10 Figures for Chapter 4

FIGURE 4.1: Share of all imports in the EU and US from China and all low wage countries

Notes: Calculated using UN Comtrade data. Low wage countries list taken from Bernard, Jensen and Schott (2006) and are defined as countries with less than 5% GDP/capita relative to the US 1972-2001.

4.11 Appendix for Chapter 4

APPENDIX 4.A: TRAPPED FACTOR MODELS: THEORY AND MEASUREMENT

A Theory of Trapped Factors and Innovation

We formulate a simple model to examine the impact in the North of a removal of trade barriers against the South (see Bloom, Terry, Romer and Van Reenen, 2012 for details). We assume that factors of production can be used to produce current goods or be used to innovate (losing a period of production). The basic idea is that there are some factors of production that are partially "trapped" due to sunk costs. With a low wage country trade shock, the opportunity cost of using these factors in innovating new goods falls as demand for the old product has been reduced, so the factors may be redeployed in innovating rather than continuing to produce the old good. As a simple example, if skilled workers are no longer used to make a low-tech product but are partly trapped within firms (for example due to firm specific human capital) they will be cheaper to deploy in designing and building a new high-tech product.

To fix ideas, consider a high wage home economy endowed with unskilled workers (*U*) who can only produce old goods and earn wage *w*, and skilled workers (*S*, who have a productivity level $\underline{\theta}$ higher than unskilled workers, *U*) who can spend their time either producing or innovating. In period 0 all workers produce a competitive generic good. In period 1, skilled workers can form partnerships of size Γ if they choose to innovate. When innovating skilled workers lose a period of production but (i) they earn some profits while the product is on patent and (ii) after a period their firm-specific productivity increases through learning by doing to $\bar{\theta} > \underline{\theta}$. If the present discounted value of innovating is Π, skilled workers will innovate in period 1 if $\underline{\theta} w \Gamma < \Pi$ before they have acquired their specific skills. After innovating and learning by doing, the opportunity cost of innovating rises to $\bar{\theta} w \Gamma$, so they will cease to innovate if $\Pi < \bar{\theta} w \Gamma$. This is because the profits from innovating are less than the opportunity cost of ceasing to produce the old good. It follows that the condition to be in a stationary equilibrium is:

$$\underline{\theta} w \Gamma < \Pi < \bar{\theta} w \Gamma$$

We consider an economy in a stationary equilibrium that has a "China shock": a trade liberalization with a low wage country on a measure of old goods that makes them unprofitable to produce but does not change the value of innovating (as by assumption China is not able to innovate in the new goods). The "China shock" thus lowers the opportunity cost (from $\bar{\theta} w \Gamma$ to $\underline{\theta} w \Gamma$) of the workers with firm-specific skills engaging in innovation. Thus, so long as the equilibrium condition holds, the China shock will induce more innovation.

The model has two further predictions we can take to the data. First, integration with another high wage country will not have the same innovation effect, as workers in these countries are paid a similar wage and old products can still be profitably produced. This is consistent with our results as we do not find any effect of imports from high wage countries on innovation. In terms of welfare, this model suggests a new benefit in addition to the usual consumer benefits of lower prices when integrating with China if there is underinvestment in R&D. [126] A second prediction is that the magnitude of the impact of innovation is increasing in the size of the trapped factor (indexed by $\bar{\theta} > \underline{\theta}$). If we allow this to be heterogeneous across industries or products, then it follows that there will be a larger impact of the trade liberalization for those

[126] In the model, underinvestment occurs even in the absence of knowledge externalities because the differentiated good sector is produced under monopolistic competition. The monopoly distortion implies that rents from innovation are lower than the total surplus as consumer surplus is ignored in the private innovation decision. An R&D subsidy would be the first-best policy, but in the absence of sufficiently high subsidies trade is a second best policy that could help close the gap between private and social rates of return to innovation.

sectors/firms with a higher level of trapped factors. We test this by interacting the Chinese import effect with proxies for the trapped factor. We turn next to how do we may measure these.

Measured inter-industry wage premia as an indicator of Trapped Factors

The model directly implies that a measure of trapped factors is the degree of product-specific skills which should be reflected in higher wages. So in principle this could be measured by the firm-specific component of a wage equation after all other general human capital and labour market shocks are controlled for. Unfortunately, such matched worker-firm data with human capital characteristics is available for only a tiny sub-sample of our firms. Consequently we turn to the more standard route of estimating a Mincerian wage equation with a full set of three digit industry dummy variables (e.g. Krueger and Summers, 1988, who use more aggregated industry dummies). The coefficients on the industry dummies are the inter-industry wage premia which, in our context will be a measure of the product-specific human capital. To do this we use pooled cross sections from the UK Labor Force Survey (LFS), the European equivalent of the US CPS (although unlike the CPS there is luckily no top-coding of the wage data). We chose the UK because it is the least regulated labor market in Europe – we did not want the results to be strongly affected by unions, minimum wages and other country-specific labor market institutions. The UK also has good publicly available quality hourly wage data on representative cross sections 1996-2008 covering our sample period.

To be precise we estimate OLS ln(hourly wage) equations of the form:

$$\ln w_{lt} = \psi' x_{lt} + \sum_j \vartheta_j IND_j + \xi_{lt}$$

Where w_{lt} is the hourly wage of worker l in year t (1996,….,2007), *IND* is a dummy for each of the $j = 1,\ldots,J$ three digit industries, and x_{lt} is a vector of wage equation controls that includes education level (dummies for four levels), a quadratic in age (for labour market experience), gender, year effects and regional dummies. We also checked the results were robust to conditioning on a sample of male workers only and to dropping all years prior to 2001 (when China joined the WTO). For example, when using the pre-2001 LFS sample, the results for Table 8, column(3) were 2.807(0.960) for the interaction term and 0.394 (0.068) for the linear industry wage premia term.

In the LFS the wage information is asked in the first (of five) quarters a respondent is interviewed and in the last quarter. We use both quarters and all years between 1996 and 2007 giving us a total of 107,622 observations for manufacturing industries (which is the sample we use).

Measured TFP as an indicator of Trapped Factors

In the trapped factor model, some firms have firm-specific inputs that generate higher productivity (e.g. workers with firm specific skills). Normalize $\underline{\theta}$=1 so that the labor services, *L*, are $U_i + \bar{\theta}_i S_i$. Assume that we can write the production function as Cobb-Douglas so

$$y_i = a + \alpha_l l_i + \alpha_k k_i + \alpha_m m_i$$

Where Y=value added, L = labor services, K = capital services and lower case letters denote logarithms so y = lnY, etc. “True” TFP is therefore:

$$TFP_i = y_i - \alpha_l l_i - \alpha_k k_i - \alpha_m m_i = \ln Y_i - \alpha_l \ln(U_i + \bar{\theta}_i S_i) - \alpha_k \ln K_i - \alpha_m \ln M_i$$

Denote measured TFP as MFP where

$$MFP_i = \ln Y_i - \alpha_l \ln L_i - \alpha_k \ln K_i - \alpha_m \ln M_i = \ln Y_i - \alpha_l \ln(U_i + S_i) - \alpha_k \ln K_i - \alpha_m \ln M_i$$

Consequently measured TFP will be equal to true TFP plus a term that depends on the importance of the trapped factors:

$$MFP_i = TFP_i + \alpha_l \ln\left(\frac{U_i + \bar{\theta}_i S_i}{U_i + S_i}\right)$$

If there are no trapped factors then $\bar{\theta}_i = 1$ and measured and true TFP are the same. Firms which have more trapped factors, $\bar{\theta}_i > 1$, however, will have a higher level of *MFP*. Thus the level of MFP for a firm is correlated with the magnitude of the trapped factors. This can be generalised to any factor which is trapped. If TFP is calculated based on the shares of the untrapped factor, then MFP will be correlated with the size of the trapped factor.

The advantage of using MFP instead of industry wage premia as a measure of trapped factors is that (i) this is firm specific time varying measure rather than an industry specific non-time varying measure and (ii) it is more general than simply being related to trapped factors in labor. The disadvantage of this measure is that it is more indirect. For example, if there is heterogeneity in the effect of trade by true TFP, then the coefficient on the interaction effect of trade and MFP in the patent equation reflects this effect as well.

APPENDIX B: DATA

Datasources

The basic data sources are described in the text, but we give some more details here.

Amadeus Accounting Data - The Amadeus data is provided by the private sector company Bureau Van Dijk, BVD. It has panel data on all European countries' company accounts. It includes private and publicly listed incorporated firms (i.e. not sole proprietors or partnerships). The accounting data includes variables such as employment, sales, capital, profits, materials and wage bills. The data goes back to the late 1970s for some countries, but is only comprehensive across a range of countries since the mid-1990s. We use successive cohorts of the Amadeus DVDs because although all firms are meant to be kept for at least 10 years after exiting, this rule is sometimes violated. Although Amadeus is a reasonably comprehensive list of names (and locations, industries and owners) for the 12 countries we study, the accounting items listed are limited by national regulations. For example, profits are generally required to be disclosed by all firms, but employment is sometimes a voluntary item for smaller firms; some countries (e.g. France) insist on wider disclosure of data than others (e.g. Germany) and disclosure is greater for public firms than private firms. In the regressions (such as the patents regressions), we consider results without and with these accounting items to check against selection bias. In terms of cleaning the accounts variables are winsorized at the 1st and 99th percentiles. The profit/sales variable winsorized between -1 and +1. Amadeus tracks the number of four digit "primary" and "secondary" four digit sectors that a firm operates in. We give primary sectors a two-third weight and secondary sectors a one third weight (results are robust to alternative weighting schemes) and weight equally within these groups (Amadeus does not report the split of sales across the four digit sectors). Using these firm-specific imports measures gives similar results to allocating all firms to their primary four-digit sector (compare Tables 1 and 5B).

EPO Patents Counts - Patents data is obtained from the electronic files of the European Patent Office (EPO) that began in 1978. We take all the patents that were granted to firms and examine the assignee names (see Belenzon and Berkovitz, 2010). We match these to the population of European firms using Amadeus (i.e. we do not insist that we have any accounting data in Amadeus when doing the matching to obtain the maximum match). The matching procedure was based on names and location. Patents are dated by application year.

In principle, a firm in Amadeus that was not matched to the EPO has taken out no patents. Nevertheless, there is a concern that we may have missed out some of the patenting activity by some firms due to the matching procedure as we were quite conservative (we only made a match if we were quite sure that the patent did belong to the Amadeus firm). We consider a narrow sample where we only keep firms if they have made at least one patent since 1978 ("patenters sample") and a wider sample where we assume that firms who we could not match really did zero patents. The analysis of patenting equations (e.g. Table 1) just uses the patenters' sample (there is no variation in the non-patenters sample) whereas the employment and survival equations (Table 3) consider both samples (see Table A2 for descriptive statistics across different samples). When constructing *PATSTOCK*, the patent stock, (e.g. Table 3) we follow Blundell et al (1999) and estimate these by perpetual inventory methods using a depreciation (δ^p) rate of 15%. $PATSTOCK_{it} = PAT_{it} + (1-\delta^P)PATSTOCK_{it-1}$ where PAT_{it} is the count of patents of firm *i* in period *t* and δ^p=0.15.

EPO Patent Citations- The EPO also provides all the citations to these patents from later EPO patents, so we use this to gauge how important a patent was (all else equal, a more highly cited patent is deemed to be more important). This is used in Table A3.

R&D - Research and Development expenditure are taken from BVD's Osiris database. These are publicly listed firms (so a sub-set of Amadeus) but Osiris contains a wider range of accounting items that Amadeus does not include, such as R&D. R&D is not a mandatory item to disclose for all publicly listed firms in Europe. Typically only the larger firms are required to disclose this item, although rules are stricter in some countries than others (e.g. in the UK under the SSAP(13) Revised accounting standard disclosure of R&D is mandatory for medium sized and larger firms). For the industry level values used in Table 5B we use the OECD's ANBERD database which is more comprehensive than Osiris as it covers all firms in a country-industry pair. We use BERD, the research and development expenditure conducted by all businesses within an industry. The ANBER data is available on a consistent basis since 1987 at a broadly two-digit industry level.

Information Technology (IT) - The IT data is drawn from an entirely different database as companies do not report IT spending except rarely as a voluntary item. Harte Hanks (HH) is a private sector company that surveys establishments in order to obtain indicators of their use of hardware, software and IT personnel. The unit of observation is a "site" which in manufacturing is a plant, so it is more disaggregated than the Amadeus data that is firm level. HH surveys plants in firms with 100 employees or more. This covers about 80% of European manufacturing employees, but obviously misses employees in smaller firms (unlike Amadeus). Each plant has an in-depth report including numbers of PCs and laptops, which we use to construct our basic computers measure. There is also information on a number of items of software such as ERP, Databases and Groupware that we use in Table A4. We have data from Harte Hanks between 2000 and 2007.

Survival - For the HH data we have a plant level measure of survival which is based on exit from the economy (i.e. *SURVIVAL* = 0 only if the plant shuts down). For the Amadeus firm-based measure we have a firm-based measure that includes both exit to bankruptcy and exit to takeover and merger (the data cannot distinguish between these types of exit).

Management data - Our management data was collected in 5 waves between 2002 and 2010. We interviewed plant managers in medium sized manufacturing firms across twenty countries (see Bloom and Van Reenen, 2007, 2010). We used a "double blind" survey tool to assess management quality across 18 questions in the areas of shopfloor operations, monitoring, targets and incentives. Each individual question is scored on a scale of 1 (worst score) to 5 (best practice) and we average across all 18 questions by firm-year observation for an overall management quality score. Each wave has a cross sectional and a panel element, with the panel element growing larger over time. There were 778 interviews in the first major cross-sectional wave in 2004 and 2,311 interviews in the last wave in 2010. Hence, we have a larger sample of data towards the end of the period with a relatively short time-span per firm. To merge the management data into the yearly trade data we linearly interpolated scores between survey waves for the same firm. So, for example, a firm which received a management score of 3.0 in 2008 and 3.2 in 2010 would have an interpolated score of 3.1 in 2009. Since we cluster by industry-country in all regressions, the t-statistics are not inflated through this interpolation procedure. The reason for this interpolation is that it increases the size of the data sample we can use in long differences – for example a firm surveyed in 2006, 2008 and 2010 could not be used in a three-year long-difference estimation without interpolating. Because the industry definitions in the management panel are not available at the four-digit level for all countries, we match industry trade data in at the three digit by country level.

UN Comtrade - Our study uses data at the HS6 product level taken from the UN Comtrade online database. We use standard concordances of HS6-SIC4 (e.g. Pierce and Schott (2010)) to aggregate to the four-digit industry level. We calculate a "value share" measure of import penetration as per Bernard, Jensen and Schott (2006) where the value of Chinese imports for a given country-SIC4 cell is normalized by the value of total world imports flowing into the same cell.

Eurostat Prodcom Production database - In Table A7 we use measures of four-digit industry-level production to normalize our imports variable. This measure of domestic production is constructed from the Eurostat Prodcom dataset. Prodcom is an eight-digit product level database of production across EU members. The first four digits of the Prodcom product code correspond to the four-digit NACE classification system. We construct a concordance between the NACE codes and US SIC, after which we aggregate the production figures to the SIC4 level. In the final step of constructing the data we compare the estimated value of production by industry-country cell to the value of exports reported in Comtrade for the same industry-country cell. In cases where the value of exports exceeds the estimated value of production from Prodcom we use the exports number as our lower bound estimate of production. This problem occurs in a limited number of cases and is due to confidentiality restrictions on the reporting of data for small industry-country cells in Prodcom.

Eurostat Producer Prices - We take two-digit industry producer prices from the online Eurostat Structural Business Statistics (SBS) database. The year 2005 is set as the base year for the price index. In some cases the data extends back to 1990 with good coverage after 1996. The SBS database reports prices in NACE codes and we concord these to the US SIC2 level to facilitate the merging in of other variables. We assemble this information for the 12 countries we focus on across our study.

Offshoring measure - This is calculated by using the US BEA input-output matrix, matched up to the Comtrade at the four-digit industry level. The offshoring variable for each industry-year is the estimated share of Chinese imported inputs in total imported inputs estimated on a similar basis to Feenstra and Hanson (1999). For each industry *j* we consider the input-output weights, $w_{jj'}$, between *j* and every other *j'* industry (note $w_{jj'}$ is from the US so the weights do not vary by country and time period). We define offshoring to China as $OFFSHORE_{jkt}^{CH} = \sum_{j'} w_{jj'} IMP_{j'kt}^{CH}$. We also considered the share of total imported inputs (from China and all other countries) in all inputs (or all costs) like the original Feenstra and Hansen paper (this replaces $IMP_{j'kt}^{CH}$ with $IMP_{j'kt}$ in the offshoring definition). However, as with our analysis of total imports in the final goods market in Table 6, it is the Chinese share (reflecting low wage country imported inputs) that is the dominant explanatory factor.

Trade weighted exchange rate IV - Following Bertrand (2004) we define each four-digit industries' exchange rate as the country-weighted exchange rate based on the source of imports in the industry. For example, an industry in Switzerland, which imported 50% from France and 50% from the UK, would have an industry-weighted exchange rate of 50% from the Euro and 50% from Sterling. This weight is held fixed by industry in the base year, but the country-specific exchange rates fluctuate every year.

Constructing industry codes

The HH plant level data (used for IT) only has a single four-digit SIC code, but this does change between years so can be used to look at product switching. Note that in Table 11 the sample conditions on firms staying within the manufacturing sector if a switch occurs i.e. plants that switch to the service sector are dropped from the sample (approximately 11% of plants switch industry according to this criterion). The Osiris data (used for R&D) only has a primary three-digit code. The Amadeus data (used for the patents, TFP and employment equations) has multiple four-digit industry codes which we can exploit to construct a weighted average of industry level imports variable to compare to the single industry code. Unfortunately, the industry data is not updated regularly so it is not reliable as a time series measure of industry switching.

The analysis of patents and TFP in the baseline specifications is based on these multiple four-digit industries. The underlying data is based on successive cross-sections of "primary" and "secondary" industry codes taken from Amadeus. We extract four cross-sections for each available year between 2003-2006. Our set of cross-sections begins in 2003 because Amadeus only began reporting primary and secondary codes separately at this point in time.

For the multiple industry import measure we use the 2003 cross-section to define a baseline set of primary and secondary four-digit industry codes for each firm. We assign a two-thirds weight to the primary codes and one-third to the secondary codes to calculate a multiple four-digit measure of import penetration (the results are not sensitive to the exact weights used). We take the arithmetic mean within sets of primary and secondary codes, that is, we weight industries equally. We follow the same procedure for calculating import penetration for the alternative normalizations presented in Tables 7 and

A7. In our data the median firm had one primary industry, the average firm 1.93 and the maximum was 10, only 19% of firms reported any secondary industry code with a mean of 2.68 and maximum of 11).

When calculating a single industry code we use the most commonly occurring four-digit code pooling across all years in the dataset. We take the lowest four-digit industry value in cases where codes occur an equal number of times. Results using this method are shown in Table 5.

APPENDIX 4.C: PRODUCTION FUNCTION ESTIMATION

The Basic Olley-Pakes Approach

Consider the basic value-added production function as:

$$y_{it} = \alpha_l l_{it} + \alpha_k k_{it} + \gamma X_{jt} + \omega_{it} + \eta_{it} \tag{C1}$$

The efficiency term, ω_{it}, is the unobserved productivity state that will be correlated with both output and the variable input decision, and η_{it} is an independent and identically distributed (i.i.d) error term. X_{jt} are the other exogenous variables in the model which are common to all firms in the industry, such as the level of quotas against Chinese goods. Assume that the capital stock is predetermined and current investment (which will react to productivity shocks) takes one period before it becomes productive, that is:

$$K_{it} = I_{t-1} + (1-\delta)K_{it-1}$$

It can be shown that the investment policy functions are monotonic in capital and the unobserved productivity state.

$$i_{it} = i(k_{it}, \omega_{it}, X_{jt}) \tag{C2}$$

The investment policy rule, therefore, can be inverted to express ω_{it} as a function of investment and capital, $\omega_t(i_{it}, k_{it}, X_{it})$. The first stage of the OP algorithm uses this invertibility result to re-express the production function as:

$$\begin{aligned} y_{it} &= \alpha_l l_{it} + \alpha_k k_{it} + \gamma X_{it} + \omega_t(i_{it}, k_{it}, X_{jt}) + \eta_{it} \\ &= \alpha_l l_{it} + \phi(i_{it}, k_{it}, X_{jt}) + \eta_{it} \end{aligned} \tag{C3}$$

where $\phi(i_{it}, k_{it}, X_{jt}) = \phi_t = \omega_t(i_{it}, k_{it}, X_{jt}) + \alpha_k k_{it} + \gamma X_{jt}$. We approximate this function with a series estimator and use this first stage to get estimates of the coefficients on the variable inputs. The second stage of the OP algorithm is:

$$y_{it} - \alpha_l l_{it} = \alpha_k k_{it} + \gamma X_{jt} + \omega_{it} + \eta_{it} \tag{C4}$$

Note that the expectation of productivity, conditional on the previous period's information set (denoted Ω_{t-1}) is:

$$\omega_{it} \mid (\Omega_{it-1}, S_{it} = 1) = E[\omega_{it} \mid \omega_{it-1}, S_{it} = 1] + \xi_{it} \tag{C5}$$

where $S_{it}=1$ indicates that the firm has chosen not to shut down. We model the selection stage by assuming that the firm will continue to operate so long as its productivity is greater than a threshold productivity, ϖ_{it}. So the exit rule is $S_{it}=1$ if $\omega_{it}\geq\varpi_{it}$, otherwise $S_{it}=0$. Taking expectations:

$$E[\omega_{it}\mid(\Omega_{it-1},S_{it}=1)]=E[\omega_{it}\mid\omega_{it-1},S_{it}=1]=E[\omega_{it}\mid\omega_{it-1},\omega_{it-1}\geq\varpi(k_{it},X_{it})]=g(\omega_{it-1},\varpi(k_{it},X_{it}))$$

We do not know ϖ_{it}, but we can try to control for it using information on observed exit.

$$\Pr(S_{it}=1\mid\Omega_{it-1})=\Pr(\omega_{it-1}\geq\varpi(k_{it},X_{it})\mid\Omega_{it-1})=\Pr(\omega_{it-1},\varpi(k_{it},X_{it}))$$

We can write the last equality as a non-parametric function of lagged observables:

$$\Pr(S_{it}=1\mid\Omega_{it-1})=P_{it}=s(i_{t-1},k_{it-1},X_{it-1})$$

So returning to the second stage coefficient of interest:

$$E(y_{it}-\alpha_l l_{it}\mid\Omega_{t-1})=\alpha_k k_{it}+\gamma X_{jt}+g(\omega_{it-1},\varpi_{it})=\alpha_k k_{it}+\gamma X_{jt}+h(\omega_{it-1},P_{it})$$

Including the shocks we have:

$$y_{it}-\alpha_l l_{it}=\alpha_k k_{it}+\gamma X_{jt}+g(\omega_{it-1},\varpi_{it})+\varsigma_{it}+\eta_{it}=\alpha_k k_{it}+\gamma X_{jt}+h(\varphi_{it-1}-\beta_k k_{it-1}-\gamma X_{jt-1},P_{it})+\varsigma_{it}+\eta_{it} \quad \text{(C6)}$$

Where $\varsigma_{it}+\eta_{it}$ are now uncorrelated with k_{it}. Since we already have estimates of the ϕ_{t-1} function and the P_{it} this amounts to estimating by Non-Linear Least Squares. We now have all the relevant parameters of the production function.

Our Implementation of Olley and Pakes

We used panel data from AMADEUS to estimate production functions between 1996 and 2006. Only four European countries had good coverage of all the factor inputs needed to estimate production function – France, Italy, Spain and Sweden. The main problem is that most countries do not insist on disclosure of both materials and capital for all unlisted private firms.

Following de Loecker (2011) we use a modified version of the Olley and Pakes (1996) approach. We allow endogeneity of the variable factor inputs (labor, capital and materials) using a control function approach and for selection through a non-parametric correction (in practice we use a second order series estimator). In addition we allow the trade variables to enter directly into the non-parametric controls for endogeneity and selectivity. As de Loecker (201) emphasizes, it is important to allow for this in order for the estimator to be consistent when the trade environment changes. We allow for imperfect competition by assuming that there is monopolistic competition which implies that the coefficients on the production function are a mix between the technological parameters and a mark-up term. The latter is identified from the coefficient on an additional control for industry output in the production function. Since some firms produce in multiple industries the relevant output term is firm-specific depending on the firm's distribution across industries. We exploit the fact that Amadeus reports the number of primary and secondary four-digit industries a firm operates in to construct this.

We do not have information on skill groups at the firm level so we also estimated TFP using the wage bill (rather than employment) as a measure of labor services, *L*. The idea is that wages reflect the different skill levels of workers in the firm, so multiplying the quantity of labor by its wage reflects the full value of labor services.

We use this method to obtain an estimate of the pure technological parameters and construct an estimate of TFP which is the variable used in the main part of the paper. We checked that the results were robust to many alternative assumptions such as estimating each parameter separately for each two-digit and country pair and by three-digit industry; allowing for higher order terms in the series approximation. Results were robust to these changes.

APPENDIX 4.D: THE TEXTILE AND CLOTHING QUOTA RELAXATION AS A QUASI-EXPERIMENT

History of trade barriers in textiles and quotas and the WTO

In 2005 restrictions on the fourth (and final) set of products regulated by the Agreement on Tariffs and Clothing (ATC) were removed. The ATC was the successor to the Multi-Fiber Agreement (MFA). The removal of quotas under the ATC came in four stages (1995, 1998, 2002 and 2005) but because China only joined the WTO in December 2001, it did not benefit initially from the first two stages. China enjoyed a substantial fall in these quotas between the end of 2001 (when it joined the WTO) and 2005 (when the ATC quotas were essentially all removed). Brambilla et al (2010) describe how there was a huge jump in Chinese exports into textiles and clothing to the US during this period (e.g. 7 percentage points increase in China's share of all US imports in 2005-2006 alone). China's increase was substantially larger than other countries not just because it joined the WTO but also because the existing quotas seemed to bite more heavily on China as indicated by the higher "fill rates" of Chinese quotas. This seemed to be because under the ATC/MFA Chinese quotas were increased more slowly over time than those in other countries.

Although formally quotas fell to zero in 2005, for 22 product groups domestic industries successfully lobbied for some "safeguards" which were re-introduced after 2005. Nevertheless, these were much lower than the pre-existing quotas. As noted in the test we only use beginning of period quotas (in 2000) to avoid the problem that post 2005 quotas are endogenous to the growth of Chinese imports. The quota policy is EU wide. It is reported in the form of the SIGL (System for the Management of Licenses for Textile Imports) database that is available online at http://trade.ec.europa.eu/sigl/choice.html. This database is classified according to 172 grouped quota categories defined by the EU. However, these categories are closely based on HS6 products so we are able to map them into the US four-digit industry classification. In addition, we added in quotas on footwear and tableware products as described in the WTO's articles of accession articles of accession for China, available at http://www.wto.org/english/thewto_e/acc_e/completeacc_e.htm. These included a selection of footwear products in the 6401-6404 HS4 categories as well as tableware products in the HS 6911-6912 range.

Construction of the Instrument

For each four-digit industry we calculated the proportion of product categories that were covered by a quota in each year (data on the amount produced in each industry is not available so we use a simple mean proportion of products). For the five-year change in imports 2005 to 2000 in the technology equations we use the quota variable in 2000 immediately prior to China's WTO entry. Specifically, this proportion represents the share of all quota-affected HS6 products in the four-digit industry (we weight each HS6 in an industry by its 2000 import value). The idea is that the market expected at this point all the quotas to be lifted. Using the actual change renders similar results, but there is a concern that the quotas remaining in 2006 are endogenous as they were the result of lobbying by the effected sectors. The "fill rates" (the proportion of actual imports divided by the quota) for most quotas were close to 100% for China in the late 1990s implying that these constraints were binding[127]. This also limits anticipation effects, although to the extent that they exist this will make it harder for us to identify a first stage. The products upon which the quotas were set were determined in the 1950s to 1970s (Spinanger, 1999) which makes them likely to be exogenous to any post 2000 actual (or anticipated) shocks. To be specific, in the regression sample of Table 2 Panel A we use all four digit US sectors in SIC4 two-digit industries 22, 23, 28, 30 and three-digit industries 314 and 326. We show that the results are robust to dropping all four-digit industries within this group with zero quotas against China in 2000 and dropping the tableware and footwear quotas.

Anticipation of China's Accession to WTO? Problems and solutions

Even if there was an unanticipated component of the China shock, since firms knew China was going to join the WTO in 2001 does this invalidate the instrument? In a stylized way one can imagine two points at which firms will react. There is an "announcement" effect on the day China's accession is determined (Costantini and Melitz, 2008, emphasis this) and an "accession" effect when China joins. For the instrument to have power in the first stage (which it does empirically), all we need is that there was some uncertainty over the effects of the accession or that firms do not fully adjust between announcement date and accession. The instrument could still be invalid, however, because the increase in technological investments (or imports) prior to accession as a result of announcement may be correlated with post-accession investments (or imports).

Formally, say the true model has the dynamic form (say because of adjustment costs)

$$\Delta \ln TECH_{it} = \lambda \Delta \ln TECH_{it-5} + \beta_1 \Delta IMP_{it}^{CH} + \beta_2 \Delta IMP_{it-5}^{CH} + u_{it} \tag{D1}$$

where *TECH* is one of our technological indicators (we use a lag of five years to be consistent with the five year differences). However, say we estimate our basic empirical model as:

$$\Delta \ln TECH_{it} = \alpha \Delta IMP_{it}^{CH} + \upsilon_{it} \tag{D2}$$

Even under the assumption that our quota instrument, Z_{it-5}, satisfies the exclusion restriction $E(Z_{it-5}u_{it}) = 0$ an IV estimation of equation (D2) using the quota instrument will be inconsistent if quotas are correlated with $\Delta \ln TECH_{it-5}$ or ΔIMP_{it-5}^{CH} due to anticipation effects. Under this assumption $E(Z_{it-5}\upsilon_{it}) \neq 0$ because υ_{it} includes the omitted lagged technology and imports variables ($\Delta \ln TECH_{it-5}$ and ΔIMP_{it-5}^{CH}). Of course, since we are estimating in long differences, it may be that $\lambda = \beta_2 = 0$ in equation (D1) so IV estimation of equation (D2) will consistently estimate α even in the presence of partial anticipation effects.

There are several ways to tackle the potential problem of anticipation effects. A direct method is to explicitly estimate the dynamic model of equation (D1). This is demanding in data terms, because we need to use firms where we observe ten full years of technology data. There are too few firms to accomplish this task for IT and patents. However, it is possible to do this for TFP and we reported the results in Table A5. We found that our results were completely robust to using the alternative dynamic specification of equation (D2).

[127] We attempted to use the fill rates in order to get a more refined measure of the instrument, but it had no additional power due to the uniformly high fill rates. Similarly, dropping all industries whose fill rates were less than 80% made no difference to the results for the same reason.

A second approach is to examine directly whether quotas are correlated with pre-WTO Accession trends in technology or Chinese imports. In our data there is a positive but small and statistically insignificant correlation between pre-WTO growth of technology (and Chinese imports) and quotas. Turning first to technical change if we regress the growth of TFP 1996-2000 (we do not have data pre-1996) on the quota instrument the coefficient (standard error) on quotas is 0.024(0.031). After China joined the WTO the five year difference 2000-2005 is 0.190(0.021) and the four year difference is 0.122(0.018). Similarly the standard reduced form for patent growth 2000-2005 has a coefficient on quotas of 0.264(0.088) whereas the regression of the pre-WTO growth of patents 1996-2000 on the quota IV has a coefficient (standard error) of 0.096(0.177).

We turn to pre-policy import trends in Table A11. We use the country by four-digit industry level information over the 1990-2007 period (we do not need technical change measures for this experiment so can use a longer period) and show regressions where the five year growth in Chinese imports is the dependent variable. Column (1) includes simply the quota (in 2000), and the positive coefficient on this variable indicates that industries where quotas were high had faster growth in Chinese imports throughout the period. Column (2) then interacts the quota variable with a policy dummy equal to one after China joined the WTO in 2001. The coefficient on this interaction is large and statistically significant, whereas the linear term on quota is small and statistically insignificant. The coefficients suggest that prior to China's joining the WTO in 2001 industries with high quotas (i.e. where all products where subject to some form of quota restriction) had 0.002 percentage point growth a year in Chinese imports (this is consistent with increases in the "fill rates" of quotas over this period as China grew). After China joined the WTO and quotas were relaxed this rose by 0.84 (= 4.2/5) percentage points per annum, a substantial amount. Column (3) includes an even more rigorous specification where we include industry dummies, allowing for industry trends over time. The coefficient on the policy-based instrumental variable remains significant with a similar magnitude of 0.04, implying that there was an increase in the Chinese growth trend post 2001.

Examples of patents taken out in the textiles and apparel industry

While the textiles and apparel sectors are relatively low tech, they were still responsible for 21,638 European patents in our sample period. These cover innovations such as new materials (for example the water resistant fabric described below), new designs (for example the more flexible ski-boat fastener described below) and new products (for example the design of an orthotic sock designed to aid ankle movement described below). Many more examples can be obtained simply by searching on the EPO web-site[128] for an appropriate textile or fabric term such as "shirt", "handbag" or "cotton".

Patent EP1335063, taken out by a German firm for a "*Water vapor permeable, water-resistant composite material*"

This is for a waterproof fabric used in, for example, protective clothing. The fabric prevents liquid water from penetrating through while at the same time permitting moisture vapor such as perspiration to pass out through the article, similar to Gore-Tex. The article has two main layers: a microporous hydrophobic outer layer which permits the passage of moisture vapor but resists penetration by liquid water; and a hyrophilic inner layer permitting the transfer of moisture vapor but preventing surface tension lowering agents such as those contained in perspiration and/or body oils from reaching the hydrophobic layer.

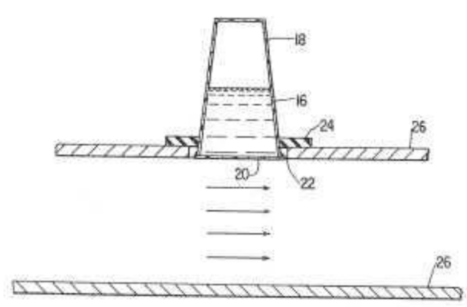

[128] http://worldwide.espacenet.com/quickSearch?locale=en_EP

Patent: EP2082659, taken out by an Italian firm for a "*Fastening device for sports footwear*"

This patent is for a more flexible in-line skate or ski boot fastener. This allows adjustment of the angle of forward inclination of the skater's leg, the circular direction of the boots and also the overall tightness of the fastening. The fastener can also include a forward inclination pressure adjusting mechanism to adjust the pressure applied to the skater's leg by the boot when the skater moves forwardly. This boot fastener can be used for a variety of purposes, with the key one being in-line skating (roller-blading), ski and snowboarding boots, but also other semi-hard sports boots and work boots.

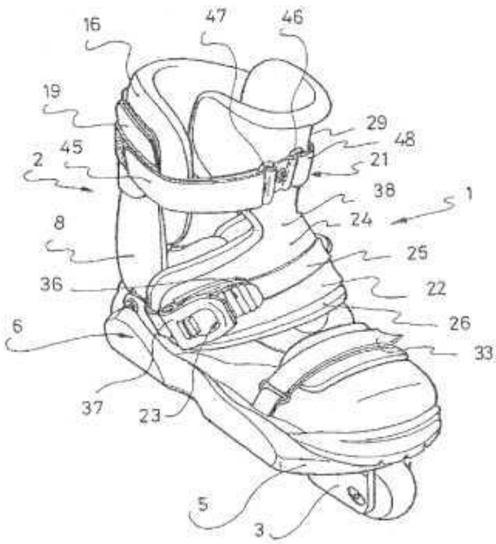

Patent: EP1626686, taken out by a UK firm for an "*Orthotic sock*"

This product provides an ankle-foot orthosis (a product to support the ankle) that comprises: an elastic structure formed of contiguous first and second tubular members, with the second tubular member set at an angle to the first tubular member to define, at least in use, a generally L-shaped cavity configured to accept and fit closely about the foot and ankle of a patient; and a rib which is permanently bonded to a region of the structure which overlies the dorsum of the patient's foot in use, with this being formed of a flexible material that has a resilience appropriate for resisting the particular degree of plantarflexion experienced by the patient.

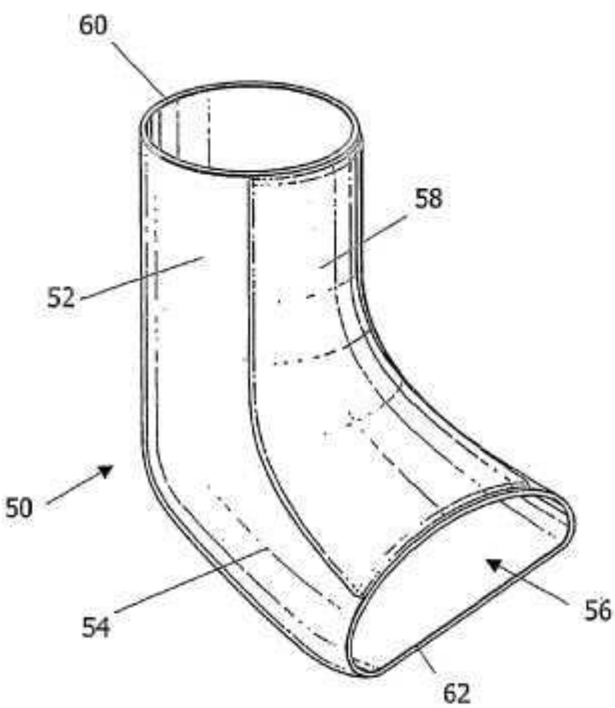

APPENDIX E: CALCULATING MAGNITUDES

The magnitudes in Table 4 attempt to quantify the potential contribution of Chinese imports to the overall increase in patents per worker, IT per worker and TFP among European manufacturing firms. Our basic approach to these calculations stems from the literature on productivity decompositions, for example, Bailey, Hulten and Campbell (1992). To explain this approach start by denoting P_t as a generic index of technology, for example aggregate patents, computers per person, or TFP. We can summarize the change in this aggregate technology index between time t and time 0 as:

$$\Delta P_t = \sum_{i=1}^{N} s_{it} p_{ijt} - \sum_{i=1}^{N} s_{i0} p_{ij0} \tag{E1}$$

where P_t, the aggregate level of the technology index, is given as a function of individual firms' technology levels (p_{ijt}) weighted by their employment shares (s_{it}), where s_{it} = firm employment divided by total employment in manufacturing. We will use patents per employee as our example, but the calculation is the same for IT per worker or TFP. This aggregate change can be decomposed into a variety of within and reallocation terms as follows:

$$\begin{aligned}\Delta P_t = &\sum_{i=1}^{N} s_{i0}(p_{ijt} - p_{ijo}) + \sum_{i=1}^{N}(s_{it} - s_{i0}) p_{ij0} + \sum_{i=1}^{N}(s_{it} - s_{i0})(p_{ijt} - p_{ij0}) \\ &- \sum_{i \in exit} s_{it}^{exit}(p_{ij0}^{exit} - \overline{p}_{jo}) + \sum_{i \in entrant} s_{it}^{entrant}(p_{ijt}^{entrant} - \overline{p}_{jt})\end{aligned} \tag{E2}$$

where $\overline{p}_{jt}$ is the average technology level of all firms in industry j year t, p_{ij0}^{exit} is the technology level of an exiter, $p_{ijt}^{entrant}$ is the technology level of an entrant and the summations are over the N firms in the economy. In this breakdown in equation (E2) the first term is the *within* effect (the increase in technology holding firm size constant), the second term is the *between* component (the increase in technology from shifting employment from low-tech to high-tech firms), the third term is the *cross* effect (the correlation of the increase in technology within firms and their change in employment share)[129]. The fourth term is the *exit* component (the impact of the relative technology level of exiting firms versus incumbent firms) and the final term the *entry* component (the impact of technology level of entering firms versus incumbent firms). As noted in the text, we cannot directly model entrants because we do not observe their lagged technology levels. In the paper we can indirectly examine the effect of entry by comparing the industry level estimates to the four components we can identify.

[129] Following the convention, we will aggregate the cross effect with the between effect when presenting results, but in practice this makes little difference as the cross-term is always small.

We have explicitly modeled the main components of these terms in our econometric models of equations (1) - (4) in the main text. Given our estimates of these in Tables 1, 2 and 3 we can create predicted values for these observable components arising from the increase in Chinese imports (ΔP_t^{China}) as follows:

$$\Delta P_t^{China} = \sum_{i=1}^{N} s_{i0}\alpha^{PAT}\Delta IMP_j + \sum_{i=1}^{N}(s_{it}^{between} - s_{i0})p_{ij0} + \sum_{i=1}^{N}(s_{it}^{between} - s_{i0})\alpha^{PAT}\Delta IMP_j - \sum_{i\in exit} s_{it}^{exit}(p_{ij0}^{exit} - \bar{p}_{jo}) \qquad \text{(E3)}$$

where α^{PAT} is the coefficient on Chinese imports in equation (1) in the main text. In Table 1 column (1) this is 0.321. $s_{it}^{between}$ is the predicted share of employment for incumbent firms and s_{it}^{entry} is the predicted share of employment in exiting firms (defined below),

$$s_{it}^{between} = \frac{N_{i0}(1+\alpha^{N}\Delta IMP_j + \gamma^{NP}\Delta IMP_j p_{ij0})}{\sum_{i=1}^{N} N_{i0}(1+\alpha^{N}\Delta IMP_j + \gamma^{NP}\Delta IMP_j p_{ij0})} \qquad \text{(E4)}$$

Where α^{N} is the coefficient on Chinese imports in the employment growth equation (equation (3) in the main text) and γ^{NP} the coefficient on Chinese imports interacted with the technology variable. The values of these are -0.352 and 1.546 respectively from column (2) in Table 3 Panel A. N_{i0} is employment in the firm[130].

$$s_{it}^{exit} = \frac{N_{i0}(1-\alpha^{S}\Delta IMP_j - \gamma^{SP}\Delta IMP_j p_{ij0})}{\sum_{i=1}^{N} N_{i0}(1-\alpha^{S}\Delta IMP_j - \gamma^{SP}\Delta IMP_j p_{ij0})} \qquad \text{(E5)}$$

Where α^{S} is the coefficient on Chinese imports in the survival equation (equation (4) in the main text) and γ^{SP} is the coefficient on Chinese imports interacted with the technology variable. In column (2) of Table 3 Panel B these are -0.122 and 0.391. Note that in equation (E5) there is a negative sign before the coefficients because we estimate survival equations econometrically whereas the decomposition uses exit.

Given these equations we can then quantify the share of technical change predicted to arise from Chinese imports as the ratio $\Delta P_t^{China}/\Delta P_t$. Similarly, we can identify the contribution of each component. To calculate ΔP_t for IT intensity we calculate the total increase in technology in our sample firms, that is, the change in the weighted mean we observe in our sample. For patents we cannot use our sample because of: (i) delays in the provision of firms accounts (we match to firm accounts and some of these are not available yet for 2005/06 due to reporting delays) and (ii) processing delays at the European Patent Office since we only use granted patents (dated by their year of application). As a result we use instead the aggregate growth of the US Patent Office (which provides long-run total patent numbers) over the proceeding 10 years

130 Note that we re-weight employment throughout the calculations so that the regression sample is representative of the entire population of Amadeus firms. This avoids any differences in data sampling or matching rates affecting the aggregate calculations.

(1996-2005), which is 2.2%. This growth rate of total patents is stable over long-run periods, for example being 2.4% over the proceeding 20 years period of 1986 to 2005.[131] Similarly, for TFP we use 2% as our measure of the growth rate of TFP growth in manufacturing in recent years.[132]

APPENDIX F: DYNAMIC SELECTION BIAS

The dynamic selection problem

Consider the representation of our baseline equations (we ignore other variables for notational simplicity) as:

$$y_{it} = \alpha z_{it} + u_{it} + \eta_i + \varepsilon_{it} \tag{F1}$$

$$s_{it} = \pi w_{it} + u_{it} + \upsilon_{it} \tag{F2}$$

where y_{it} is the technology outcome (e.g. IT/N) of interest for firm *i* at time *t* (we suppress the industry-country *jk*-subscripts), z_{it} is Chinese imports and s_{it} = 1 if the firm is operating at time *t* and zero otherwise. We assume z_{it} is exogenous, but endogeneity can easily be allowed for by using the quota instrument, for example. Assume that the idiosyncratic error terms, ε_{it} and υ_{it} are i.i.d. and the vector w_{it} includes z_{it} .

The selection problem arises from the fact that u_{it} can affect survival as well as being correlated with z_{it} .To see this consider the differenced form of equation (F1) and take expectations conditional on surviving:

$$E(\Delta y_{it} \mid \Delta z_{it}, s_{it} = 1) = \alpha + E(\Delta u_{it} \mid \Delta z_{it}, s_{it} = 1) \tag{F3}$$

The potential bias arises from the $E(\Delta u_{it} \mid \Delta z_{it}, s_{it} = 1)$ term. Under the assumption that we have instruments for Chinese imports (or they are exogenous) this simplifies to $E(\Delta u_{it} \mid s_{it} = 1)$. If the selection was solely in terms of the fixed effect, η_i or captured by the observables w_{it} , then this expectation would be zero and our estimate of the effect of trade would be unbiased, so "static selection" is not a problem. The concern is that there is "dynamic selection" on technology shocks, Δu_{it} , so $E(\Delta u_{it} \mid s_{it} = 1) \neq 0$.

To see the dynamic selection problem in our context consider two industries A and B, one (industry A) has an increase in Chinese imports (e.g. from a relaxation of quotas) and the other (B) has not. Now consider the reaction to this shock of two identical firms who both have had the same negative productivity shock unrelated to China. If the firm in industry A is more likely to exit (as life will get harder in the future) then it will appear that within firm productivity growth improves in

[131] The data goes back to 1986 on aggregate USPTO patents and comes from http://www.uspto.gov/go/taf/cbcby.htm. The EPO does not have this long-run of time series aggregate patents data since it was only founded in 1977 and was not widely accepted (over European national patent offices) until the late 1980s making the time series unreliable prior to the 1990s.

[132] The growth rate of European multifactor productivity growth 1995-2008 was 1.9% per annum according to Conference Board (http://www.conference-board.org/economics/downloads/Summary_Statistics_2010.pdf, taken from Table 12 for the EU-12).

industry A, even if nothing else changes. Although there is a genuine increase in productivity in industry A as more of the low productivity firms are "cleansed" by Chinese competition, we attribute too much of this to the within firm component.

One strategy for dealing with this problem is to consider "instruments" for survival i.e. variables that effect the probability of survival that do not affect the technology shock. This is the standard Heckman (1979) selection equation where we would include selection correction terms generated from equation (F2) augmented to equation (F3). It is difficult to think of such exclusion restrictions in our context, however, that could enter w_{it} but be excluded from z_{it} [133]. Instead we take two alternative approaches: (i) placing a lower bound on the selection bias and (ii) adopting a non-parametric control function approach to control for the bias.

Bounding the Selection Bias

A recent literature in econometrics emphasises that even when point identification is not feasible, it may be possible to achieve set identification. In our context, this means that we may be able to place a lower bound on the effect of Chinese imports on technology. Following Manski (1994), Manski and Pepper (2000) and Blundell et al (2007) we consider the "worst case bounds", i.e. what could be the lowest effect of imports if selection effects were severe. What helps out in our application is that there is a finite lower support at zero for the distribution of patents and IT. If the firm had survived it could never have less than zero patents or zero computers. In this case we can impute that all the exiting firms would have performed zero patents and lost all their computers had they survived. Any positive effect remaining from α will be the "worst case" bounds.

Control function approach for selection

The worst-case bounds approach is infeasible for TFP as it is a continuous variable without finite support. One approach would be to use less conservative bounds (e.g. assume that the exiters were all from the lowest decile of the TFP distribution). These approaches need some rather arbitrary cut-off rule so instead we the same control function approach suggested by Olley and Pakes (1996) to add a non-parametric term in the propensity score based on observed exiters when estimating the production function. This is described above in Appendix C (e.g. see equation C6)

[133] Some possibilities based on alternative (usually strong) dynamic assumptions include Kyriazidou (1997), Honore and Kyriazidou (2000) or Wooldridge (1995).

TABLE A1: CHINA'S SHARE OF GLOBAL IMPORTS – TOP TEN INDUSTRIES, 1999-2007

Top Ten Industries in 1999 (by China's import share)		**China's Share of all Imports** IMP^{CH}		
Industry Description	Industry Code	1999	2007	Change 2007-1999
Dolls and Stuffed Toys	3942	0.817	0.859	+0.042
Drapery, Hardware and Window Blinds	2591	0.527	0.574	0.047
Rubber and Plastics Footwear	3021	0.532	0.618	0.086
Leather Gloves and Mittens	3151	0.517	0.574	0.057
Women's Handbags and Purses	3171	0.470	0.517	0.047
Manufacturing NEC	3999	0.458	0.521	0.064
Games, Toys and Children's Vehicles	3944	0.434	0.765	0.331
Luggage	3161	0.432	0.680	0.248
Personal Leather Goods	3172	0.416	0.432	0.016
Apparel and other Finished Fabric Products	2386	0.415	0.418	0.003
All Industries		0.057	0.124	0.068
(standard-deviation)		(0.102)	(0.152)	(0.089)

Notes: Calculated using product-level UN Comtrade data aggregated to four-digit US SIC codes. There are 430 four-digit industries in our dataset. China's share of all imports IMP^{CH}_{1999} total world imports. Countries include Austria, Denmark, Finland, France, Germany, Ireland, Italy, Norway, Spain, Sweden, Switzerland, the UK and the US. the Manufacturing industries (not elsewhere classified) includes many miscellaneous goods such as hairdressing equipment, tobacco pipes, cigarette holders, artificial flower arrangements, and amusement or arcade machines.

TABLE A2: DESCRIPTIVE STATISTICS

Variable	Mean	Stan. Dev.	Median	Minimum	Maximum
Patenters sample - Firms with at least one EPO patent since 1978					
Number of Patents (per firm-year)	1.022	10.40	0	0	882
Employment	739.5	6,526.7	100	1	463,561
Number of Observations	30,277				
IT sample (Harte-Hanks)					
Number of Employees	248.3	566.1	140	1	50,000
IT Intensity (computers per worker)	0.580	0.385	0.398	0.05	2.00
Industry switchers (% plants switching four-digit sector in five year period)	0.112	0.316			
Number of Observations	37,500				
R&D sample (Osiris)					
R&D/Sales ratio	0.152	0.888	0.034	0.001	17.3
Employment	17,230	46,422	2054	4	464,841
Number of Observations	1,626				
TFP sample (Amadeus)					
Employment	79.4	333.9	30	10	84,252
Number of Firms (in TFP sample)	89,369				
Number of Observations	292,167				
Management sample					
Management score	3.11	0.58	3.14	1.11	4.89
Employment	716	902	350	100	5,000
Number of firms	1576				
Number of observations	3,607				
Employment sample (Amadeus)					
Number of Patents (per firm-year)	0.019	5.741	0	0	882
Employment	99.95	1,504.9	17	1	372,056
Number of Observations	581,474				
Survival sample (Amadeus)					
Number of Patents	0.049	2.80	0	0	830
Employment	97.8	2,751.7	14	2	1,469,840
Number of Observations	490,095				

Notes: Standard deviations in parentheses. Samples are based on those used to run regressions, so we condition on having non-missing values over a five-year period for the relevant variable. “Patenters sample” are those firms who have at least one patent in the European Patent Office (EPO) since 1978. Employment sample is based on Amadeus (again firms have to have reported employment over a five-year period as this is the dependent variable in the regressions. IT sample is HH. IT intensity is computers per worker. R&D sample is from Osiris (publicly listed firms). TFP sample is Amadeus firms in France, Italy, Spain and Sweden. Management sample covers all firms in France, Germany, Italy, Ireland, Sweden and the UK with multiple management interviews.

TABLE A3: NO FALLS IN CITATIONS PER PATENTS BECAUSE OF CHINESE IMPORTS

Dependent Variable:	**Δln(CITES)**	**Δ ln(CITES/PATENTS)**
Change in Chinese Imports	0.118	0.009
ΔIMP_{jk}^{CH}	(0.081)	(0.029)
Number of industry-country clusters	1,578	1,578
Number of Firms	8,480	8,480
Observations	30,277	30,277

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country by four-digit industry pair in parentheses. Estimation by five-year differences. ΔIMP^{CH} represents the 5-year difference in Chinese imports as a fraction of total imports in a four-digit industry by country pair. All specifications include country-year fixed effects. 12 Countries. Sample period is 1996 to 2006. *Δ(CITES)* is defined as the change in *ln(1+CITES)* where *CITES* = count of citations and *Δ(CITES/PATENT)* is defined as the change in *ln[(1+CITES)/(1+PAT)* where *PAT* = count of patents.

TABLE A4: ALTERNATIVE IT ADOPTION MEASURES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	ΔERP (ENTERPRISE RESOURCE PLANNING)			ΔDATABASE			ΔGROUPWARE		
Change in Chinese Imports	0.040			0.002			0.249***		
ΔIMP_{jk}^{CH}	(0.034)			(0.070)			(0.083)		
Highest Quintile for ΔIMP_{jk}^{CH}		0.013***			0.020**			0.034**	
		(0.005)			(0.010)			(0.014)	
2nd Highest Quintile of ΔIMP_{jk}^{CH}		0.006			0.030***			0.021	
		(0.005)			(0.010)			(0.013)	
3rd Highest Quintile for ΔIMP_{jk}^{CH}		0.014***			0.043***			-0.008	
		(0.005)			(0.010)			(0.013)	
4th Highest Quintile for ΔIMP_{jk}^{CH}		0.010**			0.024***			-0.018	
		(0.005)			(0.011)			(0.013)	
Lowest Quintile for ΔIMP_{jk}^{CH}			-0.011***			-0.028**			-0.000
			(0.004)			(0.009)			(0.001)
Number of Observations	24,741	24,741	24,741	24,741	24,741	24,741	24,741	24,741	24,741

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation by OLS with standard errors (clustered by country by four-digit industry pair) in parentheses. There are 2,728 distinct country by industry pairs. . Quintiles represent bands of establishments ordered from highest (5) to the lowest (1) in terms of their change in Chinese Imports, that is, quintiles of ΔIMP^{CH} . 12 Countries. All regressions have site-type controls, employment growth and country by year dummies

TABLE A5: CHECKING PRE-POLICY TFP TRENDS

	(1)	(2)	(3)	(4)
Dependent Variable:	**ΔTFP**	**ΔTFP**	**ΔTFP**	**ΔTFP**
Estimating Method:	**IV**	**IV**	**IV**	**IV**
Δ Chinese Imports$_t$	1.897**	1.491***	1.608***	1.635***
	(0.806)	(0.264)	(0.410)	(0.313)
ΔTFP$_{t-5}$			-0.211***	0.378***
			(0.024)	(0.063)
ΔChinese Imports$_{t-5}$			-0.531	-0.450
			(0.602)	(0.423)
Endogenous right-hand side variables	Chinese Imports	Chinese Imports	Chinese Imports	Chinese Imports, ΔTFP(t-5)
Number of units	55,791	3,107	3,107	3,107
Number of clusters	187	126	126	126
Observations	55,791	3,107	3,107	3,107

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by four-digit industry in parentheses. These are estimates from the textile and apparel industries following Table 2 Panel A. Five-year differences covering the period 1999-2005. Estimation by five-year differences. Quota removal is based on EU SIGL data and defined as the (value weighted) proportion of HS6 products in the four-digit industry that were covered by a quota restriction on China in 1999 (prior to China's WTO accession) that were planned to be removed by 2005 (see the Appendix D for details). In columns (1)-(3) we use quota removal to instrument Chinese imports. In column (4) we also use TFP_{t-10} as an instrument for ΔTFP_{t-5}. 4 Countries.

TABLE A6: DYNAMICS OF THE EFFECT OF CHINA ON PATENTS AND EMPLOYMENT

PANEL A: PATENTS, Δln(PATENTS)	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
5-year lag of Change in Chinese Imports	0.328***					
ΔIMP^{CH}_{t-5}	(0.110)					
4-year lag of Change in Chinese Imports		0.394***				
ΔIMP^{CH}_{t-4}		(0.110)				
3-year lag of Change in Chinese Imports			0.402***			
ΔIMP^{CH}_{t-3}			(0.120)			
2-year lag of Change in Chinese Imports				0.333***		
ΔIMP^{CH}_{t-2}				(0.113)		
1-year lag of Change in Chinese Imports					0.314***	
ΔIMP^{CH}_{t-1}					(0.102)	
Contemporaneous change in Chinese Imports						0.321***
ΔIMP^{CH}_{t}						(0.102)
Number of country-industry pairs	1,578	1,578	1,578	1,578	1,578	1,578
Number of Firms	8,480	8,480	8,480	8,480	8,480	8,480
Observations	30,277	30,277	30,277	30,277	30,277	30,277

PANEL B: EMPLOYMENT, Δln(N)	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
5-year lag of Change in Chinese Imports	-0.188					
ΔIMP^{CH}_{t-5}	(0.140)					
4-year lag of Change in Chinese Imports		-0.241*				
ΔIMP^{CH}_{t-4}		(0.139)				
3-year lag of Change in Chinese Imports			-0.306**			
ΔIMP^{CH}_{t-3}			(0.155)			
2-year lag of Change in Chinese Imports				-0.275*		
ΔIMP^{CH}_{t-2}				(0.160)		
1-year lag of Change in Chinese Imports					-0.285**	
ΔIMP^{CH}_{t-1}					(0.143)	
Contemporaneous change in Chinese Imports						-0.309**
ΔIMP^{CH}_{t}						(0.138)
Number of country-industry pairs	1,464	1,464	1,464	1,464	1,464	1,464
Number of Firms	7,030	7,030	7,030	7,030	7,030	7,030
Observations	22,938	22,938	22,938	22,938	22,938	22,938

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country by four-digit industry pair in parentheses. All columns estimated as 5-year differences ΔIMP^{CH}_{t-l} represents the 5-year change in Chinese imports (where l = lag length). 12 Countries. Sample period is 1996 to 2005.

TABLE A7: ALTERNATIVE MEASURES OF THE CHANGE IN CHINESE IMPORTS

PANEL A: CHINESE IMPORTS NORMALIZED BY DOMESTIC PRODUCTION

Dependent Variable:	(1) Δln(PATENTS)	(2) Δln(IT/N)	(3) Δln(TFP)	(4) Δln(N)	(5) SURVIVAL
Change in Chinese Imports (over production)	0.142***	0.053**	0.065***	-0.232***	-0.103***
$\Delta\left(M_{jk}^{China}/D_{jk}\right)$	(0.048)	(0.024)	(0.020)	(0.033)	(0.017)
Change in Chinese imports*ln(Patent stock per worker at t-5)				0.507	0.456***
$\Delta\left(M_{jk}^{China}/D_{jk}\right)$ *(PATSTOCK/N)$_{t-5}$				(0.431)	(0.111)
ln(Patent stock per worker at t-5)				0.503***	0.041***
(PATSTOCK/N)$_{t-5}$				(0.054)	(0.009)
Number of Units	8,474	20,106	89,369	189,309	488,270
Number of industry-country clusters	1,575	2,480	1,210	3,115	3,335
Observations	30,221	31,820	293,167	579,818	488,270

PANEL B: CHINESE IMPORTS NORMALIZED BY APPARENT CONSUMPTION

Dependent Variable:	(1) Δln(PATENTS)	(2) Δln(IT/N)	(3) Δln(TFP)	(4) Δln(N)	(5) SURVIVAL
Change Chinese Imports (over apparent consumption)	0.349***	0.169*	0.045**	-0.477***	-0.203***
$\Delta\left(M_{jk}^{China}/C_{jk}\right)$	(0.122)	(0.089)	(0.019)	(0.078)	(0.034)
Change in Chinese imports*ln(Patent stock per worker at t-5)				1.385	0.476***
$\Delta\left(M_{jk}^{China}/C_{jk}\right)$ *(PATSTOCK/N)$_{t-5}$*				(1.238)	(0.187)
ln(Patent stock per worker at t-5)				0.490***	0.041***
(PATSTOCK/N)$_{t-5}$				(0.078)	(0.009)
Number of Units	8,474	19,793	89,369	189,309	488,270
Number of industry-country clusters	1,575	2,406	1,210	3,115	3,335
Observations	30,221	31,225	293,167	579,818	488,270

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country by four-digit industry pair in parentheses. $\Delta\left(M_{jk}^{China}/D_{jk}\right)$ represents the 5-year difference Chinese Imports normalized by domestic production (D). $\Delta\left(M_{jk}^{China}/C_{jk}\right)$ is the 5-year difference in Chinese imports normalized by apparent consumption (C). Apparent consumption defined as Production - Exports + Imports (C=D-X+M). Variables D and C is from Eurostat's Prodcom database with full details given in the Data Appendix. Quintile 1 is a dummy variable for firms in the lowest quintile of IT intensity in the baseline year. Note that Switzerland is not included because it does not report production data to Eurostat's Prodcom database. Sample period is 2000 to 2007 for the IT equation and 1996-2005 for patents equations. Column (2) controls for the growth in employment.

TABLE A8: OFFSHORING TO CHINA – FULL RESULTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable:	**Δln(PATENTS)**	**Δln(IT/N)**	**Δln(TFP)**	**Δln(N)**	**Δln(N)**	**Δln(N)**	**SURVIVAL**	**SURVIVAL**	**SURVIVAL**
Measure of Lagged TECH:				**Patent stock**	**IT**	**TFP**	**Patent Stock**	**IT**	**TFP**
ΔIMP_{jk}^{CH}	0.313***	0.279***	0.189***	-0.392***	-0.269***	-0.374***	-0.090	-0.110	-0.172**
	(0.100)	(0.080)	(0.082)	(0.145)	(0.105)	(0.103)	(0.060)	(0.079)	(0.074)
$\Delta IMP_{jk}^{CH} * TECH_{t-5}$				0.142*	-0.362**	0.679	0.339**	0.071	0.053
				(0.086)	(0.168)	(0.477)	(0.167)	(0.138)	(0.075)
$\Delta OFFSHORE_{jk}^{CH}$	0.173	1.685***	1.396***	-1.643	-2.802***	-0.227	-0.500	-1.546***	-0.533**
	(0.822)	(0.517)	(0.504)	(1.202)	(0.682)	(0.544)	(0.316)	(0.550)	(0.223)
$\Delta OFFSHORE_{jk}^{CH} * TECH_{t-5}$				1.064	1.406	4.874**	1.950	1.315**	0.568
				(0.70)	(1.111)	(2.181)	(2.030)	(0.710)	(0.411)
$TECH_{t-5}$				-0.012	0.219***	0.231***	0.016	-0.125	-0.007
				(0.008)	(0.013)	(0.019)	(0.018)	(0.008)	(0.005)
Number of units	8,480	22,957	89,369	6,335	22,957	89,369	1,647	2,863	1,294
Number of industry-country clusters	1,578	2,816	1,210	1,375	2,816	1,210	7,985	28,624	268,335
Observations	30,277	37,500	292,167	19,844	37,500	292,167	7,985	28,624	268,335

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation by OLS with standard errors (clustered by country by four-digit industry pair) in parentheses. ΔIMP^{CH} represents the 5-year difference in Chinese imports as a fraction of total imports in a four-digit industry by country pair. The variable $\Delta OFFSHORE$ is the 5-yerar change in Chinese imports in source industries, defined following Feenstra and Hansen (1999) – see Appendix B. Countries in all columns except for TFP models which is for four countries. Columns(1)-(3) repeat the results reported in Table 10. Columns (4)-(6) repeat the analysis of employment changes in Table 3 Panel A but also include the control for offshoring (and its interaction with lagged technology). Columns (7)-(9) repeat the analysis of survival (conducted in Table 3, Panel B) with a control for offshoring (and its interaction with lagged technology). All columns include country by year effects. 12 countries (except in column (3), (6) and (9) which are four countries).

TABLE A9: MAGNITUDES ALLOWING FOR OFFSHORING

All Figures are as a % of the total increase over the period 2000-2007

PANEL A: Increase in Patents per employee attributable to Chinese imports

Period	Within	Between	Exit	Total
Product Market	4.6	7.6	2.0	14.2
Product market + Offshoring	5.1	8.0	1.4	14.5

PANEL B: Increase in IT per employee attributable to Chinese imports

Period	Within	Between	Exit	Total
Product Market	9.8	3.1	1.2	14.1
Product market + Offshoring	20.9	5.3	3.3	29.5

PANEL C: Increase in Total Factor Productivity attributable to Chinese imports

Period	Within	Between	Exit	Total
Product Market	10.1	4.4	0.3	14.7
Product market + Offshoring	24.6	7.6	0.8	33.0

Notes: Panel A reports the share of aggregate patents per worker accounted for by China, Panel B the increase in IT per worker and Panel C the increase in total factor productivity. In each panel the first row ("Product Market") simply reports the same results following methodology in Appendix E implemented in Table 4 (the results differ slightly from Table 4 because we only use the single industry version of Chinese imports as in Table 5 Panel B as the multiple industry version is not available for offshoring). We then extend the methodology to allow for offshoring to China. All underlying regression specifications are extended to allow for offshoring to China. The full specifications of the within firm (same as Table 10), between and exit specifications are those in Table A8. We multiply the relevant coefficients by the observed Chinese import share growth to generate a predicted change in IT/Employee, Patents/Employee and TFP between 2000 to 2007 inclusive. The lower row in each panel ("Product Market + Offshoring") decomposes the total change (final column) into within, between and exit effects for the combined product market and "offshoring elements.

TABLE A10: EXPORTS TO CHINA

Dependent Variable:	(1) Δln(PATENTS)	(2) Δln(IT/N)	(3) ΔTFP
Change in Chinese Imports ΔIMP_{jk}^{CH}	0.322***	0.361***	0.254***
	(0.102)	(0.076)	(0.072)
Change in Exports to China $\Delta\left(X_{jk}^{China} / X_{jk}^{World}\right)$	-0.243	0.028	-0.125
	(0.200)	(0.118)	(0.126)
Number of Units	8,480	22,957	89,369
Number of Industry-country clusters	1,578	2,816	1,210
Number of Observations	30,277	37,500	292,167

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by country by four-digit industry in parentheses.12 Countries except column (3) where there are four countries. "Number of units" represents the number of firms in all columns except (2) where it is plants. 12 countries except in column (3) where it is four countries.

TABLE A11: THE QUOTA INSTRUMENT IS UNCORRELATED WITH THE GROWTH IN CHINESE IMPORTS PRIOR TO THE ACCESSION TO THE WTO

Dependent Variable	(1) ΔIMP^{CH}	(2) ΔIMP^{CH}	(3) ΔIMP^{CH}
Quota Removal*Post WTO		0.042***	0.039***
		(0.010)	(0.010)
Quota Removal	0.036***	0.009	
	(0.008)	(0.008)	
Country by Year Effects	Yes	Yes	Yes
Country by industry trends	No	No	Yes
Number of clusters	84	84	84
Observations	11,138	11,138	11,138

Notes: *** denotes 1% significance; ** denotes 5% significance; * denotes 10% significance. Estimation is by OLS with standard errors clustered by four-digit industry pair in parentheses. This data is a four-digit industry by country panel between 1990 and 2007. Sample is the textiles and clothing industries only. The dependent variable is the five-year difference in Chinese import share. Quota removal is the height of the quota in the four-digit industry in 2000 prior to China joining the WTO. "Post WTO" is a dummy equal to unity after 2001 (and zero before). 12 countries.

CHAPTER 5: REAGAN'S INNOVATION DIVIDEND? TECHNOLOGICAL IMPACTS OF THE 1980S US DEFENSE BUILD-UP.

.

Abstract

US government spending since World War II has been characterized by large investments in defense related goods, services and R&D. In turn, this means that the Department of Defense (DoD) has had a large role in funding corporate innovation in the US. This paper looks at the impact of military procurement spending on corporate innovation among publicly traded firms for the period 1966-2003. The study utilizes a major database of detailed, historical procurement contracts for all Department of Defense (DoD) prime contracts since 1966. Product-level spending shifts – chiefly centered around the Reagan defense build-up of the 1980s – are used as a source of exogenous variation in firm-level procurement receipts. Estimates indicate that defense procurement has a positive absolute impact on patenting and R&D investment, with an elasticity of approximately 0.07 across both measures of innovation. In terms of magnitudes, the contribution of defense procurement to innovation peaked during the early Reagan build-up, accounting for 11.4% of the total change in patenting intensity and 6.5% for R&D. This compares to a defense sector share in output of around 4%. The later defense cutbacks under Bush Senior and Clinton then curbed the growth in technological intensity by around 2%.

JEL Codes: 030, 031, 038, 043.

Keywords: Induced Innovation, Patenting, R&D, Defense Spending.

5.1 Introduction

Technological development is central to many debates on economic policy. At the macroeconomic level it underpins economy-wide productivity, so governments invest considerable resources to influence the development of new technologies. Surprisingly, academic research on the determinants of the rate and direction of technical change makes up only a small sub-field of economics. Schmookler's (1966) contribution outlined an agenda for studying technical change that has been followed up heavily in empirical research that has appeared over the last 10-15 years. This work has included research on areas such as medical innovation (Acemoglu and Linn 2004; Finkelstein 2004); the energy sector (Newell, Jaffe and Stavins 1999; Popp 2002); and the impact of low-wage import competition (Bloom, Draca and Van Reenen 2011; Le Large 2010).

However, it is government intervention itself that became a major driver of technical change in the mid-20th century. World War II introduced an era of 'organized innovation' centered on the defense sector (Mowery and Rosenberg 1991). The massive spending in R&D made during WII provided a model for ongoing investment following Vannevar Bush's (1945) policy report, *Science: The Endless Frontier*. As Figure 1 shows, real defense spending has hovered around an average level of around $300 billion per year since WWII with peaks during 'hot war' episodes (Korea, Vietnam, the Gulf Wars) as well as some 'cold war' phases (notably the Reagan-led build-up of the 1980s).

The scale of this spending is significant along a number of dimensions that make it important for the study of innovation. Firstly, defense spending represents a large fraction of government outlays – approximately 15-20% over this period, which is comparable to the budgets for health and social security (OMB 2008). Secondly, this spending has a large role in determining the level of total R&D expenditures. The NSF(2006) estimates that Department of Defense (DoD) funding accounts for around 20% of R&D expenditures in the post-war period, with a peak of around 30% in the late 1950s and early 1960s. Thirdly, the amount of money flowing into high-tech, defense-focused production dwarfs the amount spent on other prominent innovation policy tools. For example, the Federal R&D tax credit costs around $6.5 billion per year while support for basic science through the National Science Foundation figures at $7 billion (NSF 2006). By contrast, around $16 billion per year is spent on military R&D procurement alone along with another $40-50 billion in spending on high-tech products. This makes defense spending – and military procurement in particular – one of the most significant topics for the study of induced innovation in the US economy.

In this paper I assess the impact of military procurement on corporate innovation among listed US firms, who have historically been the main recipients of DoD spending. Specifically, I focus on the impacts of DoD funding on patenting and R&D investment and address two main questions. Firstly, what are the absolute and relative impacts of government defense procurement purchases on corporate patenting and R&D investment? The absolute impact is defined in terms of the elasticity between procurement and the two innovation measures while the relative impact is measured using the benchmark of an equivalent civilian-sector sales shocks. Secondly, given these impacts, what has been the historical magnitude of the defense sector's contribution to corporate innovation through the procurement channel? These questions are important not only because of the size of the procurement budget but also because of debate about the innovative qualities of military spending.

This debate has stressed the special characteristics military spending on a number of counts. On the negative side, it is argued that the effectiveness of defense-oriented production in fostering innovation is likely to be lower than comparable effort directed at the civilian economy. This argument is based on the specialized, mission-focus of defense production. For example, spending on defense R&D has emphasized the "D" component of development rather than "R" element of basic science[134]. Furthermore, defense production is thought to prioritize technical, battlefield performance as a main goal thereby sacrificing efficiency, cost-effectiveness and potential dual-use in the civilian sector. There is also a line of argument regarding the potential 'crowding out' and displacement effects of federally-funded R&D on private R&D[135]. Finally, defense production is also notorious for its general inefficiency and high costs which have inevitable consequences for productivity.

On the other side of the argument, it is possible that military procurement spending could have a positive impact on innovation. In particular, the high-tech nature of defense procurement sales could push out the 'innovation possibilities frontier' as firms commit their resources to more

---

[134] Figures from the NSF(2006) indicate that between 1956-2005 work classified as 'Development' accounted for approximately 80% of DoD sponsored R&D while 'Applied' work had a 10-15% share and 'Basic' R&D claimed less than 5%.

[135] The main crowding out channel here is the input market – federal spending could drive up the price of R&D inputs and lower the real value of spending (Goolsbee 1999). Investment displacement could also occur in cases where (for political or other reasons) the government has funded projects with high rate of private return. Hence the government could be subsidizing projects that private companies would have pursued with their own funds (David, Hall and Toole 2000)

technically ambitious projects than those demanded by the civilian sector. Cozzi and Impullitti (2010) present a growth model along these lines where the technological composition of government spending matters for innovation. Specifically, they outline a 'demand-pull' channel where a shift towards spending on high-tech goods increases the rewards for innovation and ultimately has an effect on the relative demand for skilled workers[136]. Ruttan (2006) also presents a detailed historical account that tracks the role of military funding in the development of crucial general purpose technologies such as jet passenger aviation, computing, nuclear energy and, most recently, the global positioning system (GPS) network. This historical literature suggests that the contribution of military spending to innovation could be significant but there is limited microeconomic evidence to support this at the level of individual firms.

In this paper I fill this gap and assemble a new firm-level dataset on corporate innovation and defense procurement. This involves the matching of a major historical database of all DoD prime procurement contracts between 1966-2003 to firm-level information on publicly traded firms (from COMPUSTAT) and the NBER US Patent Database. The prime contracts database reports very detailed information including characteristics such as contractor identity, type of product sold, extent of competition, location of work, and dates of action among other features. This level of detail greatly facilitates the analysis in terms of both the heterogeneity of spending and the scope for identifying exogenous shifts in firm procurement receipts.

Empirically, the potentially endogenous allocation of funding makes it important to identify such exogenous shifts in the amount of procurement contracts received by companies. The fact that changes in defense spending have historically been driven by strategic considerations assists in tracing out these spending shocks. In general, the procurement receipts of individual firms have been at the mercy of major turning points in defense spending such as the Korean war, Vietnam, the 'Reagan build-up' after the invasion of Afghanistan, the fall of the Berlin Wall, and latterly 9/11[137]. Historically, this means that some firms have been exposed to very sharp increases and decreases in DoD spending. In this paper, I identify exogenous shocks to

---

[136] Their calibrated model posits that this government policy channel has an impact on innovation which then explains between 12-15% of observed changes in wage inequality between 1976-2001. This is consistent with earlier evidence by Berman, Bound and Griliches (1994) which found that changes in defense shipments accounted for just over 15% of the increase in non-production worker employment shares in the 1979-87 period.

[137] Changes in military strategy have also shifted the composition of spending. Markusen et al (1990) document how the rise of the 'air power' doctrine shifted the structure of spending circa WWII. This was followed by further changes as ballistic missile technology became strategically important in the 1950s.

firm procurement receipts by using the lagged product specialization of firms interacted with the current level of DoD spending in the firm's hypothesized defense 'product market'.

Practically, this approach measures how the DoD's aggregate and product-level spending decisions affect the procurement receipts of individual firms. The central assumption here is that no individual firm can influence the level of DoD demand for goods and services at the product (or indeed aggregate) level. Since my sample covers the period 1966-2003[138] the major 'turning point' in aggregate spending is the 1980s Reagan build-up and subsequent cutbacks in the 1988-2000 period under Bush and Clinton. My analysis therefore pays special attention to the structure of the build-up and the impact of this overall spending policy.

The main results analyze the impact of defense-based sales relative all other non-defense or 'civilian' sales. The OLS estimates indicate the elasticity between defense procurement and both measures of innovation is around 0.07. This implies that a 10% increase in procurement contracts is associated with a 0.7% increase in patenting and company-sponsored R&D. Benchmarking this elasticity against non-defense or 'civilian' sales I find that the estimated effect for patents is twice as high as its civilian benchmark, supporting the idea that defense procurement is a source of demand for high-tech production. Furthermore, this effect also holds for cite-weighted patents indicating that the additional patents are not secured through a quality trade-off. The effects for company-sponsored R&D are broadly in line with the civilian sales benchmark but are much higher when total R&D (including external, defense-funded R&D) is considered. The proposed IV (firm-specific product market shocks) strongly predicts defense procurement and suggests elasticities higher than the OLS estimates for both patents and R&D. The effectiveness of the IV strategy is supported by an analysis of the DoD's product level spending patterns. Specifically, following the main identification assumption outlined above, there is no evidence that firms in concentrated industries are able to tilt DoD spending in their own favor.

Finally, in terms of historical magnitudes, the contribution of defense procurement to innovation peaked during the early Reagan build-up when the defense sector accounted for 11.4% of the total increasing in patenting and 6.5% of the change in R&D. This is a large effect given that the defense share of sales is around 4% for the full corporate sample. These magnitudes are calculated using an industry-level decomposition and it is notable that the majority of the patenting effect is due to within industry changes. This is consistent with the firm-level estimates

[138] Changes in the procurement reporting system after 2003 make it difficult to link up the data from 2004 onwards. See the data section for further notes.

which indicate that defense procurement spending strongly stimulates patenting. The later cutbacks under Bush Senior and Clinton then served to moderate the growth in technological intensity with the between industry shifts in defense output curbing the total trend increase by 2%.

The remainder of the paper is as follows. The next section offers more background and a discussion of data. The modeling approach is then outlined, followed by the discussion of results and conclusions.

5.2 Background

5.21 Defense Procurement Policy in the US

In this section I discuss three main features of US defense procurement policy relevant to this study. The first important feature is the scale of procurement spending in this area. Figure 2a plots the total amount of procurement spending by fiscal year from 1966 (expressed in 2003 prices). This data is derived from the historical files used for the empirical work in this paper. Spending peaks at over $220 billion in 1985 and records its lowest levels after Vietnam and in the mid-1990s prior to end of the Clinton administration. This total includes goods and services purchased from all types of contractors, including listed firms, non-listed firms and universities or other research institutes. Based on the calculations from my matched dataset the COMPUSTAT-listed firms are the biggest group of contractors in weighted dollar terms, accounting for 60-70% of all spending per fiscal year.

The second important feature of procurement policy relates to the structure of spending. Contracts in the DoD procurement database are classified according to a 4-digit product code. While there are over 1500 4-digit products across 155 2-digit groups, spending is still heavily concentrated amongst a subset of products. Figure 2b shows the procurement spending shares of the ‘Top 10’ products at the 2-digit level. The DoD’s purchases are massively capital and research intensive with R&D, electronics, ships, and missiles making up a large fraction of spending alongside regular armed forces supplies such as subsistence and fuel. This large share for the Top 10 products is a consistent feature and in later analysis I explore how product shares changed over time.

The final important background feature to DoD policy is the way that procurement spending decisions are taken. Specifically, the DoD purchases its major goods and services on both a competitive and non-competitive basis following a ‘life-cycle’ model. For example, when the DoD commissions a new weapons system it establishes a ‘technical design competition’ and solicits detailed, scientific proposals from potential contractors. Firms contest this stage vigorously because winning a design contest assures them of receiving non-competitive follow-on contracts. These follow-on contracts relate to supply and maintenance – since the firm designed the weapon system it has monopoly power over its ongoing provision. The DoD's power in this situation depends on its ability to substitute across weapons system which according to a study by Lichtenberg (1990) is limited[139]. The uncertainty inherent in high-tech defense projects (Peck and Scherer 1962) also means that costs for the DoD often increase after the competitive stage, enhancing the financial position of the locked-in contractor.

Practically, the DoD supports the design competition process by providing on ongoing R&D subsidy for firms so that they are primed to submit detailed technical bids. The effects of the DoD's overall policies are therefore felt by firms both on the demand-side (through procurement policy) and the supply-side (via the R&D subsidy).

5.22 Analytical Framework

The demand and supply-side effects of DoD policy on firm-level innovation can be understood using the general framework established by Hall, David and Toole (2000) who in turn built on the work of Howe and McFetridge (1976). This framework is illustrated in Figure 3 and contains some obvious components. Over each planning period, firms rank potential innovation projects according to their anticipated yield thereby forming a marginal rate of return (MRR) schedule. The volume of innovation investment here is denoted as R&D. The marginal cost of capital for funding these potential innovation projects is then traced out by the upward sloping MCC schedule. The upward slope of the MCC reflects the increasing cost of funds as the volume of R&D increases. The use of internal firm funds is represented by the flatter area of the MCC with external financing accounting for the upward slope. This very simple framework can be set up as:

$$MRR = f(R, Z_1) \qquad (5.1)$$

$$MCC = g(R, Z_2) \qquad (5.2)$$

with R standing for R&D expenditures while the Z_1 and Z_2 vectors capture 'shift variables' that respectively affect the range of project rates of return and the marginal costs of capital. Of course, the optimal R&D occurs where the MRR and MCC are equalized, such that $R^* = h(X, Z)$.

---

[139] Lichtenberg (1990) studies data on cost and quantity revisions at the weapon systems level and finds an elasticity of demand of 0.55.

The impact of federal procurement funds can then be nested according to potential shift factors. Hall et al (2000) summarize the Z_1 vector for the MRR according to three types of variables: those that affect technological opportunities; those influencing the state of demand for a firm over its lines of business; and finally, institutional factors that affect the appropriation of innovation outputs. By analogy the MCC shifters can be categorized as: firstly technology policy measures that directly target the cost of capital; then macroeconomic conditions that affect internal funds; bond market conditions underpinning the cost of external finance; and finally institutions affecting the availability of finance (i.e. venture capital).

Procurement contracts such as those offered by the DoD are therefore best framed as demand shifters affecting the MRR. As discussed, contracts for R&D services are frequently coupled with valuable, non-competitive follow-on contracts for the final goods designed as part of the research process. The large demand component therefore has the effect of moving the firm MRR schedule outwards. Furthermore, this is not a secular increase in demand but rather one that applies for military-related innovation investments. Following the introductory discussion, the increased emphasis on military-related investments could then have a positive or negative effect on the subsequent innovation outputs of firms (chiefly measured in this paper by patents).

This type of shift in the MRR contrasts with the effects of a direct R&D grant or subsidy. Such policies increase the effective level of internal funds and shift the MCC schedule to the right. It should be noted that there is a clear distinction in DoD policy between procurement and subsidy-based R&D funds. The DoD administers a subsidy policy known as the Independent Research and Development (IR&D) program. This program reimburses firms for the overhead costs incurred as part of non-contract work that is related to military R&D priorities broadly defined. The work is independent in the sense that the research projects involved are selected and initiated by the private company itself. The main objective of the program is to underwrite the efforts of firms in participating in technical design competitions for new projects as outlined in the previous section. The role of the IR&D program was studied in detail by Lichtenberg (1989, 1990).

Given this overall background, this paper mainly treats procurement spending as working through the demand–side of the firm innovation investment decision. While some shifts in the

MCC could still be induced by changes in procurement spending policy[140] the majority of procurement effects are likely to fall on the demand-side. However, the presence of the IR&D program complicates the analysis. This program operates in parallel to procurement spending and will have the effect of pushing the MCC outwards.

Systematic data on the IR&D subsidy is not as readily available as the information on procurement but some conclusions about its influence can be drawn. Firstly, the level of the IR&D subsidy is determined by a formula that depends in part on lagged level of procurement spending. As such, the IR&D subsidy will be correlated with firm procurement receipts, albeit with a delay[141]. Secondly, while the IR&D subsidy is valuable to firms it is still small compared to the total value of procurement. For example, Lichtenberg (1990) calculates that allowable costs under the IR&D policy were worth $3.5billion in 1986 for all contractors which represents only 2.5% of the total procurement budget. The ultimate implication of this correlation between the two policies (that is, the IR&D subsidy and procurement spending) is that the reduced form estimates I present below will be picking up some degree of shift in the MCC along with the bigger effects of procurement on the MRR.

5.3 Data

Three main datasets are used to build the long-run firm panel used in this paper: historical military procurement data from the National Archives and Records Administration (NARA); US Patent and Trademark Office (USPTO) information on patents (as compiled by Hall, Jaffe and Trajtenberg (2002) as part of their NBER project); and company accounts information from COMPUSTAT. The details of each dataset are discussed in turn.

5.31 DoD Procurement Contracts

The NARA historical files on military procurement contain all prime military contracts awarded by the Department of Defense (DoD) since the 1966 Fiscal Year (FY) and until FY 2003. After 2003, the DoD changed its procurement reporting format. It began to report its procurement

[140] For example, this could occur in cases where the signal of demand from the DoD makes firms more attractive to external sources of finance and in situations where procurement has a ‘pump priming’ effect that lowers the fixed costs of research.

[141] Lichtenberg (1989,1990) explains that the IR&D subsidy is calculated as the firm’s defense sales-to-total sales ratio multiplied by a ceiling amount for allowable R&D costs. This ceiling is determined by lagged R&D expenditures claimed under the program. In turn, this means that firms are able to claw back expenditures that exceed the ceiling in the current period because the overspend has the effect of ratcheting up the ceiling in future periods.

information as part of the highly complicated Federal Procurement Data System (FPDS). I plan to include post-2003 data from the FPDS in future iterations of the paper.

The file for each FY contains records on approximately 250,000 different contracts awarded by all DoD sub-agencies for the purchase of goods and services. The records are drawn from a standardized departmental form known as the DD 350 or more eloquently as the "Individual Contracting Action Report". The minimum reporting threshold for purchases is $10,000 for FY1966 – FY1983 and $25,000 for FY1984 onwards.

The data are exhaustive and summarize many details of each contract, such as: the names and unique identifiers of the awardees; contracting office within the DoD; types of contracts (e.g. competitive versus non-competitive); dates of action; estimated completion date; geographic location of the contractor (city, county and state); weapon system code; and importantly a 4-digit product code (known as the Federal Supply Code (FSC)). While there is some addition and deletion of products the FSC classification is consistently defined from 1966, making it feasible to define a 155 product panel across the 1966-2003 period. The NARA data are probably the most detailed historical data on government procurement available anywhere and were only released in this form in the late 2000s. As a result, research using these military procurement files is still very limited. Some examples of work that uses defense procurement data of this type includes Hines and Guthrie (2011) and Nakamura and Steinsson (2011), along with the Frank Lichtenberg's program of work in the 1980s and early 1990s (summarized in Lichtenberg 1995).

5.32 COMPUSTAT Accounts data

The COMPUSTAT database provides accounts information on stock-market listed firms, with annual information available from 1950 onwards. I extracted the raw data for all firms from 1966 onwards. In cleaning the sample, all accounting and procurement variables were winsorized at the 1^{st} and 99^{th} percentiles. The final sample reported in the regressions from Table 5.2 onwards drops all firms with fewer than four years of consecutive data. Furthermore, note that the sample used from Table 5.2 also conditions on the existence of a 10-year lag for procurement receipts and therefore begins in 1976. This is because the proposed exogenous shocks term is based on 10-year lagged product shares.

In terms of variable definitions, sales (mnemonic SALE) is used as the output measure; the net stock of property, plant and equipment (PPENT) is used for the book value of capital, and the labour input is represented by employees (EMP). The R&D capital stock is defined following the perpetual inventory method (PIM) using a 15% depreciation rate as $G_t = R_t + (1-\delta)G_{t-1}$

where R_t represents the flow of company-sponsored R&D expenditures (mnemonic XRD). Note that this is also the approach taken for the calculation of patent stocks using the USPTO data. The return on assets is define as Net Income (NI) over Current Assets (ACT). The return on sales from data on Sales, Cost of Goods Sold (COGS) and Selling and Administrative Costs (XSGA).

5.33 Measuring R&D

The measurement of R&D deserves special attention in the context of defense procurement spending. The flow of R&D expenditures reported in COMPUSTAT represents the sum of company-sponsored R&D. This follows the Securities Exchange Commission (SEC) definition of R&D as all costs incurred for research and development into new products, processes or services. Importantly, this SEC definition excludes customer and government sponsored R&D, including the R&D awarded to firms as part of defense procurement contracts[142]. Practically, this means that the COMPUSTAT measure of R&D is not picking up the total amount of R&D activities conducted at a firm. In contrast, the Science Research Statistics (SRS) branch of the NSF conducts a Survey of Industrial Research and Development (SIR&D) which surveys R&D spending according to all types of funds. This survey indicates that the company-sponsored R&D measured in COMPUSTAT represents around 80% of total R&D expenditures, with the remainder made up mainly of government sponsored R&D.

This measurement issue impinges on some of the econometric models estimated in this paper. For example, it means that the implied elasticity between R&D and patenting will be biased upwards for firms that receive large sums of defense procurement business. This is simply because the full sum of the firm's expenditure on R&D inputs is not factored into the company-sponsored measure that is given by COMPUSTAT.

The extent of this bias can be evaluated by comparing the COMPUSTAT measure to other measures of R&D (such as that reported in the SIR&D) that *do include* government-sponsored portion of expenditures. However, in lieu of access to the SIR&D survey I construct a measure of 'Company Plus' R&D that is based on the company-sponsored R&D reported in COMPUSTAT plus the value of the procurement contracts reported in the NARA files that are

[142] Other items excluded from the COMPUSTAT measure of R&D include: software-related expenses, the cost of extractive activities (ie: prospecting, drilling); routine engineering activities directed at product and process improvements; inventory royalties; and market research or testing (NSF 2006).

product coded as R&D. Results using this 'Company Plus' measure of R&D are reported alongside results for the company-sponsored only measure[143].

5.34 Patents Data - General

The final key dataset for the project is the NBER US Patents Database (Hall et al. 2002). These data were produced as part of an ongoing NBER project that processes raw USPTO patent data and matches patent assignees against the full historical set of stock-market listed firms. The data were first produced in 1999 with an update in 2006 and ongoing work to deepen the dataset.

The NBER Patents data provides the frame for the name matching exercise that I conduct across the three datasets. That is, I used the list of the assignees from the NBER database as the main source of names to be matched to the NARA procurement database. The string-based name-matching is implemented using the usual procedures outlined in work such as Hall et al (2002) and Bloom, Draca and Van Reenen (2010). The presence of Dun and Bradstreet (DNB) business numbers allows me to consolidate establishments in the procurement data to the HQ level before matching. For completeness, I match COMPUSTAT company names directly to the procurement database to capture cases where firms receive defense contracts but do not necessarily patent. Finally, I also manually match assignees and contractors in cases where high-value contracts cannot be matched using the automatic method. Final match rates are high in weighted dollar value terms. Approximately 78% of contracts by weighted value are matched to either the NBER Patents Database or COMPUSTAT. This rises to around 94% for contracts classed under R&D product codes.

5.35 Patents Data – Defining Military Patents

Given the focus of this paper, it is interesting to ask whether defense procurement has induced more innovation directed at defense-based technologies. To look at this I define a measure of 'military-intensive' patent classes. This measure is meant to represent patents produced under conditions where clear DoD interests can be inferred. The logic here is that patents falling in these technology classes are more likely to represent specialized military technologies.

---

[143] While defense-sponsored R&D represents a large fraction of all government-sponsored R&D (approximately 50% for the main sample considered in this paper), other federal departments will also contribute funds. Future iterations of this paper will match in data on non-defense Federal procurement contracts in order to complete this picture of total firm R&D expenditures.

These military technology classes are defined following two criteria. Firstly, I filter out all private company patents where the DoD holds some shared property rights. This is derived using the Government Support field within the *USPTO Patent Grant Full Text* files[144]. This field notes cases where a patent was 'made with Government Support' such that public agencies can claim some legal rights to the invention. Typically, the Government Support section of the full text files gives the name of the agency concerned (for example, the US Navy) and the procurement contract number. An example of a 'Government Interest' declaration in a 1987 Navy-sponsored missile patent by the Hughes Aircraft Company is given in Appendix A.2. As part of this work, I extract all cases of government-supported patents involving DoD agencies.

Secondly, these DoD Government Interest patents are then pooled with the DOD patents created in-house by the Army, Navy and Air Force in order to create a group of military specialized patents. The patents in this pool are then allocated according to their technology class, with these classes ranked according to the proportion of military patents falling within them. The distribution of military patents is highly concentrated with over 45% of patents falling into the top 10% of technology classes by rank. I then classify all patents falling within this top 10% of technology classes as the group of 'military specialized patents'. As discussed, the rationale for this is based on the fact that these are the technology classes most commonly associated with direct DOD involvement and therefore represent the classes that are closest to the production of pure, military specialized goods. Finally, patents falling in technology classes outside of this top 10% are then classified as non-military or civilian patents.

5.4 Modelling Approach

5.41 Basic Econometric Approach

This paper considers two main technological outcomes of interest: patenting and R&D spending. Theoretically, these two equations can be motivated using a simple framework. First, consider a factor demand for R&D inputs derived from a simple production function $Q = AG^{\delta}K^{1-\lambda}L^{\lambda}$ where Q is firm output, G is knowledge capital (measured by R&D) and K and L are the labor and capital inputs:

$$\ln G = \mu \ln Q - \sigma \ln[p^{g} / p^{x}] \qquad (5.3)$$

[144] The specific version of the files used is that available from the Google Patents bulk download facility located at: http://www.google.com/googlebooks/uspto-patents-grants-text.html.

where (p^g / p^x) is the relative price ratio between R&D and other types of input. Since patents are then produced by the firm mainly using these R&D inputs we can think of a firm-level patents production function defined as $PAT = BG^{\gamma}$ where *B* is an efficiency parameter. Taking logarithms and substituting in our expression for *G* above, we have:

$$\ln PAT = \ln B + \gamma\left[\mu \ln Q - \sigma \ln[p^g / p^x\right] \qquad (5.4)$$

Empirically, note that the price ratio terms in (3) and (4) are constant across firms and absorbed by time effects. The most important term for the models estimated in this paper is output *Q* (measured by reported firm sales) which enters directly into (3) and indirectly with a coefficient of $\gamma\mu$ into the patents equation defined by (4).

This term is important because the value of defense contracts received by a firm in a given period can be interpreted as a sales term subject to caveats regarding measurement error. This sales-based interpretation of defense contracts provides the basis for the two main reduced form technology equations I estimate in this paper. Firstly, consider a generic firm-level outcome equation as follows:

$$\ln TECH_{ijt} = \alpha_i + \beta_1 \ln D_{ij(t-1)} + \delta' X_{ijt} + sic_j * t + u_{ijt} \qquad (5.5)$$

where $TECH_{ijt}$ is a measure of innovation (either R&D or patents) observed at the level of firm *i* in industry *j* at time *t*; $\square_i$ is the firm fixed effect; X_{ijt} is a vector of explanatory variables; sic_j*t stands for industry-level time trends and u_{it} is the disturbance term. The key variable of interest is $D_{ij(t-1)}$, the amount of defense procurement dollars received by firm *i*, lagged by one-period here to avoid immediate feedback effects.

The most general issue for this type of reduced form is the potential bias on $\square_1$. The DoD is likely to award contracts to the most innovative and competitive firms in the market for defense-related goods, contributing to an upward bias on $\square_1$. Furthermore, the DoD also has an interest in acquiring and developing the latest technologies which means it could allocate its D_{ijt} funds according to areas of growing technological opportunity. That is, the DoD's spending could be targeting fields and product classes where *TECH* is already growing for exogenous scientific reasons, again contributing to an upward bias on $\square_1$. The inclusion of firm fixed effects $\square_i$ and industry trends sic_j*t terms are two steps that can be taken to deal with these issues and I discuss this more in the next sub-section.

A second important issue is the interpretation of β_1, the coefficient on the value of defense procurement contracts. Extending the interpretation of this variable as a form of 'defense sales' we can add a further term to this reduced form as follows:

$$\ln TECH_{ijt} = \alpha_i + \beta_1 \ln D_{ij(t-1)} + \beta_1 \ln C_{ij(t-1)} + \delta' X_{ijt} + sic_j * t + u_{ijt} \qquad (5.6)$$

where C_{ijt} is the amount of all non-defense or 'civilian' customer sales. Practically, the civilian sales are calculated by subtracting the value of defense procurement contracts reported in the NARA data from total reported sales given by firms in COMPUSTAT. Including civilian sales in this way helps interpretation by giving us a benchmark to judge the value of β_1. It can be expected that $\beta_2 > \beta_1$ simply due to the size of civilian sales relative to defense sales. That is, a 10% change in civilian sales is necessarily larger than a 10% change in defense sales. However, the comparison we are interested in is the effect of given changes in defense or civilian sales that are of equal size. This is simply a matter of normalizing the elasticities by the defense-civilian sales ratio which allows us to test whether the implied coefficient for β_1 (calculated as $(D/C)*\hat{\beta}_2$) is different to the estimated $\hat{\beta}_1$[145]. The implied coefficient here gives us the impact that a dollar of defense procurement sales would have on *TECH* if that dollar affected the firm in exactly the same way as any other dollar of civilian sales.

The main difficulty with this approach to interpreting the effects of defense procurement is the measurement error involved in treating procurement as a form of sales. At face-value, the procurement contracts reported in the NARA are indeed the administratively recorded 'sales' of defense goods and services made by a firm to the government. However, the translation of the procurement dollars reported in the NARA data to the language of company accounts is distorted at two points. Firstly, the procurement dollars reported in the contracts data are aggregated according to start-date such that many multi-year projects are recorded up-front[146]. Secondly, the accounting treatment of procurement receipts as they enter into company accounts is very complicated (Lichtenberg 1992; Rogerson 1992, Thomas and Tung 1992).

---

[145] To see this note that we are interested in $\partial \ln TECH / \partial D > \partial \ln TECH / \partial C$. Re-arranging this with respect to the elasticities gives us $\partial \ln TECH / \partial \ln D > (D/C)*[\partial \ln TECH / \partial \ln C]$.

[146] The NARA data does contain information on the both the start date and end date of contracts. Hence it is possible to allocate the sales from multi-year contracts according to these dates. Future iterations of this paper will conduct this exercise.

One key matter identified by Rogerson (1992) is the cost-shifting of overheads. Under certain contract structures firms can shift their overheads on non-defense projects onto the overhead claims made for defense projects. Other accounting issues also include: the treatment of procurement receipts as 'income' rather than revenue[147] ; interactions with non-defense policies such as the R&D tax credit; and accounting for sub-contracts within the NARA prime contracts. However, the main empirical implication here is that we can expect the measurement error to affect account-based variables such as R&D more severely than non-accounting variables such as patents. I return to this issue of measurement error in the discussion of results.

5.42 Analysis of Patents

Patenting is measured in terms of integer 'counts' of the number of applications made in a given year. This introduces a non-negligible number of zero observations. I deal with this firstly by applying the log '1 plus patents' normalization that is commonly adopted. However, this is an arbitrary truncation. In the interests of robustness a more formal count data model is necessary. This can be specified as:

$$P_{ijt} = \exp\left(\alpha_i^P + \beta_1^P \ln D_{ij(t-1)} + \beta_1^P \ln C_{ij(t-1)} + \delta^P X_{ijt} + sic_j^P * t + u_{ijt}^P\right) \qquad (5.7)$$

where P_{ijt} is the count of the firm *i* at time *t*, with the *P* superscripts denoting this as the patenting version of the general technology equation outlined previously in (6). Since the usual Poisson assumption that the variance and mean are equal is not valid in this context (the 'overdispersion' problem), I adopt a negative binomial specification. This is estimated following the conditional maximum likelihood approach of Hausman, Hall and Griliches (1984).

5.43 Exogenous Shifts in Firm Procurement Receipts

As discussed, the allocation of procurement funds is likely to follow some endogenous patterns. It is logical that the DoD will award contracts to firms that are already highly innovative – indeed the competitive structure of the procurement process is designed to do this (subject to price considerations). It is also plausible that the DoD may target areas of growing technological opportunity as part of its objective to build the best military equipment possible. To address these endogeniety issues I will take three steps: (i) control for firm-level unobservables with fixed effects; (ii) include a full set of 4-digit industry trends; and (iii) adopt an IV strategy based on

[147] For example, a firm can record the initial receipts for (say) a helicopter supply contract at first in income and then only as revenue when the final units are individually priced and delivered to the DoD.

exogenous shocks to firm procurement receipts. This section outlines the construction of the exogenous shocks variable.

The detailed product-level information reported in the DoD prime contracts data provides a rich setting for defining identification strategies using a 'shift-share' approach. This approach is based on taking the lagged product specialization of a firm or location and then calculating the current demand based on DoD procurement spending. Intuitively, the premise is that firms have a pre-existing specialization in types of goods purchased by the DoD. As the DoD varies its spending year by year then the size of the potential defense market for the firm changes. If the firm's shares across product groups are defined with a sufficient lag we can limit the influence of situations where firms endogenously enter into new product categories where the DoD is increasing spending. The lagged pattern of specialization is therefore designed to capture the firm's core products for sale to the DoD.

We can express this by first defining the historical product shares (here using a 10-year lag) for a firm:

$$\Phi_{il,(t-10)} = \frac{d_{il,(t-10)}}{\sum d_{il,(t-10)}} \tag{5.8}$$

where $d_{il,(t-10)}$ represents the amount of procurement dollars received by firm i in product category l at lagged time period *(t-10)*. This measures the firm or location's degree of specialization across a basket of products. The level of total product demand for firm i in the current period can then be calculated as follows:

$$d_{it}^{L} = \sum \Phi_{il,(t-10)} D_{ilt} \tag{5.9}$$

where D_{ilt} is the sum of all procurement spending by the defense department in product category l during current period t. The expression in (9) therefore measures how the department's spending patterns affect firm or location i based on a predetermined, historical specialization. A key assumption here is that no individual firm can affect the level of demand in product group l (e.g. through political lobbying). The efficacy of this assumption can be tested by studying the pattern of spending at the product group level and relating it to group characteristics such as concentration ratios, market power or political clout.

5.44 Calculating Magnitudes

The modeling approach presented up until now has focused on estimating the firm-level relationship between procurement and innovation outcomes. Aggregate magnitudes can be calculated using these firm-level parameters and nesting them alongside a decomposition of changes in technology. Defense procurement spending will have effects on both the 'between' firm or industry distribution of technology as well as the level of technology 'within' an industry or continuing firm. Berman, Bound and Griliches (1994) put forward an industry-level decomposition of skilled labor that incorporated the effects of defense-induced demand. Adapting this approach, a basic industry decomposition for changes in technology can be defined as follows:

$$\Delta Tech_t = \sum_j \Delta q_j \overline{Tech}_j + \sum_j \Delta Tech_j \bar{q}_j \qquad (5.10)$$

where $q_j = \left(Q_j / \sum Q_j\right)$ is the share of industry j output in total output across industries and $Tech_j = \left(TECH_j / Q_j\right)$ is a measure of technological intensity per industry (ie: patents or R&D per dollar of sales). The $\Delta Tech_t$ term represents the aggregate change in technological intensity summed across all industries. The first term on the right-hand side of (9) then represents the change in overall technology that is explained by between industry shifts in output shares with industry-level technological intensity held constant. The second term measures the within industry shift in technological intensity while holding output shares fixed. Berman et al (1994) modify this basic decomposition to account for three sectors operating within industries: domestic civilian production, defense-related production, and production for import and export. For current purposes I will focus only on the split between the civilian and defense production sectors. Assuming that technological outputs by sector are proportional to the sectoral share of industry output we can re-write the between component as:

$$\sum_j \Delta q_j \overline{Tech}_j = \sum_j \Delta q_j^D \overline{Tech}_j + \sum_j \Delta q_j^C \overline{Tech}_j \qquad (5.11)$$

where the D and C superscripts represent the defense and civilian production sectors within an industry. The within component can then be written in a similar fashion as:

$$\sum_j \Delta Tech_j \bar{q}_j = \sum_j \Delta Tech_j \bar{q}_j^D + \sum_j \Delta Tech_j \bar{q}_j^C \qquad (5.12)$$

with the same convention on the D and C superscripts. Note that the assumption that technological outputs are proportional to sectoral industry shares links back to the firm-level technology model outlined in equation (5.6). This earlier model tests whether the within-firm production of technology is more responsive to defense procurement sales compared to non-defense or civilian sales. If technology is indeed more responsive to defense procurement sales then this means the defense sector contribution measured in equation (5.11) is actually a lower-bound. In such a case, the defense sector contribution to industry-level technology will be larger than its proportional share in industry output.

5.5 Results

5.51 Basic Patenting and R&D Results

Table 5.1 summarizes the sample constructed from the matching of the NARA, COMPUSTAT and NBER Patents data. For this table, the data is divided into three types of firms: those who never receive any defense procurement funding; the firms who receive funding in some of the years that they are observed but not all; and the firms that receive funding in all of the years that they are observed. In this paper my primary focus is on the latter section of the sample whose characteristics are reported in column (3). This is to done to capture the effects of intensive margin changes and avoid the problems of censoring that occur at zero values of defense sales. The main sample based on column (3) is therefore a selected sample, albeit representative of the majority of defense expenditures in weighted dollar terms. Of course, it is possible to estimate selection models and estimate bounds to evaluate the influence of this sampling choice.

The data in Table 5.1 indicates that the firms selling goods and services to the DoD in all observed years are very research intensive, with more than three times as many patents as the firms in column (2), although mean R&D expenditures are comparable if we take account of the difference in sales. On average, defense procurement is equal to 4.9% of total firm sales for column (3) sample and this group of firms is responsible for 84% of all DOD procurement purchases from listed firms. Clearly then, the fact that the 'All Years' sample represents such a large share of DoD purchases mitigates the sampling choice since most of the government procurement spending is encompassed by this definition.

Table 5.2 then presents the basic results on patenting for variations of equations (3) and (4). Column (1) provides a basic specification with SIC4 fixed effects, resulting in a high coefficients of 0.320 and 0.409 for patenting and company-sponsored R&D respectively. This is unsurprising insofar that we expect DoD purchases to be associated with the largest and most innovative firms *ex ante*. The second column in Panels (A) and (B) includes firm fixed effects to account for the unobservable characteristics of firms. This reduces the coefficient on defense procurement by 75% compared to the initial specification. The third column then includes SIC4*year industry trends as an additional control since both patenting and R&D have experienced strong upward trends in recent decades, with some industries increasing at higher rates than others[148]. This results in a coefficient of around 0.07 for both outcomes. Applying this as an elasticity suggests that a 10% increase in procurement contracts is associated with a 0.7% increase in patenting and company-sponsored R&D. This is suggestive of a very high elasticity of around 1 between the defense-induced increase in R&D and the similar induced increase in patenting. To this end, column(4) of Panel B reports the result of a regression using the 'company-plus' measure of R&D outlined in section III. This measure directly includes defense-funded R&D in addition to company-sponsored R&D and the result in column (4) indicates an elasticity of around 0.70. This is closer in line with the observed elasticity between patents and R&D for this sample, which is approximately 0.67 (0.023).

Finally, note here also that the negative binomial model for patents (column (4)) yields a similar estimate to the OLS model which adopts the log '1 plus patents' normalization. Only around 15% of observations in this sample record a zero count for patents. As in Hausman, Hall and Griliches' (1984) study of patents, R&D and count data methods it is therefore not surprising that the two estimates are similar.

Table 5.3 tackles the question of whether this elasticity is large in terms of its economic context. This table estimates different versions of equation (6) which includes civilian sales as an additional term alongside defense procurement. As expected, the coefficient on civilian sales is much higher than that for defense sales. For example, the civilian sales coefficient for patenting is 0.45 in the basic column (2) specification compared to 0.060 for defense procurement. Following the discussion in section IV, this is down to the fact that a 10% change in civilian sales is necessarily larger than a 10% increase in defense procurement. The implied $\square_1$ reported in this table therefore normalizes the civilian coefficient by the ratio of defense procurement to civilian sales. This provides a benchmark for the effect that defense procurement would have if it were just another increment of sales.

The results indicate that the estimated $\square_1$ is significantly higher than the benchmark for patenting but not for company-sponsored R&D. Specifically, the estimated $\square_1$ for patenting is twice as high as the implied benchmark and the point estimate for company-sponsored R&D in panel (B), column(3) is over 15% higher than its relevant benchmark, although not significant. In contrast, the elasticity is much higher than the benchmark when the 'company-plus' measure of R&D is used (panel(B), column (4)). Importantly, this difference in results across the R&D

[148] See Hall, Jaffee and Trajtenberg (2001) for full details. For example, Computer and Communication related patents increase from around a 5% share in 1975 to over 15% by the mid-1990s. The share for the Mechanical category fell from over 25% to around 17-18% in the same time period.

measures gives us some indication of how defense procurement is impacting overall company R&D. Specifically, the results show that the *additional* effect of defense procurement comes into play once defense-funded R&D is measured. However, the fact that defense procurement and civilian sales have similar effects on the private, company-sponsored portion of R&D (panel (B), column (3)) counts as evidence against any 'crowding out' of private R&D expenditures by defense R&D. Finally, as in Table 5.2, the implied elasticity between 'company-plus' R&D and patenting is approximately 0.54, which is in range of the baseline estimate for this sample of 0.65.

To summarise, the results in Table 5.3 do strongly suggest that defense procurement sales are a source of demand for high-tech goods. The magnitude of this relationship for patenting indicates that this could be a very strong effect, with twice as many patents produced for a given dollar of defense sales compared to the same dollar of civilian sales. The effect on R&D is also twice as high as the civilian benchmark, although only when considering the expanded 'company-plus' measure of R&D. However, it is notable that (contrary to the 'crowding out' hypothesis) a dollar of defense sales is associated with at least as much company-sponsored R&D as a dollar of civilian sales.

5.52 Type and Quality of Patents

Table 5.4 takes the analysis of patents further by studying the quality and type of patents. Following the discussion in section III, the total number of patents produced by firm *i* in year *t* is divided into two groups: military specialized or 'defense' patents (ie: those patents belonging to technology classes that are the main focus of direct DoD invention); and non-military or 'civilian' patents (in practice, all patents belonging to technology classes outside of the military specialized set). Panel (A) of Table 5.4 provides evidence on the effect of sales on different types of patents, with the first column repeating the result for total patents initially reported in Table 5.3. There are two points to note about panel (A). Firstly, there is some suggestive evidence that military patents are more affected by defense sales than civilian sales. The civilian sales coefficient is lower in column (2) relative to column (3) and the defense sales coefficient in column(2) is further away from its implied benchmark value as a result. Secondly, it is clear that defense sales are strongly associated with both types of patents. This indicates that defense procurement has an effect on innovation outside of a limited set of military patent classes. Specifically, this means that increases in defense sales do not strongly favor military technologies over civilian technologies but rather there is a general effect on innovation across technology classes.

Table 5.4 also addresses the issue of patent quality by examining backward citations. This represents a measure of scientific importance based on the number of citations made to a patent subsequent to its granting. The issue of quality is important insofar that the higher patenting due to defense sales could be a function of firms moving down their investment curves for innovation, that is, producing a higher quantity but lower quality of innovations in terms of scientific importance. Table 5.4 tests this first by using a count of citations as the dependent variable in columns (4). This leads to a coefficient of 0.067 for the specification with industry trends which is comparable to the coefficient of 0.059 for patenting in Table 5.3. This estimate for citations is still well above the implied benchmark of 0.032 but there is less precision compared to the Table 5.2 result. The next two columns then distinguish between cite-weighted military and civilian patents. Again, the effect of defense sales is well above the implied benchmark for both measures, although the effect is stronger for military patents. The overall conclusion from Table 5.4 then is that defense procurement stimulates patenting significantly above the civilian benchmark with a minimal trade-off in terms of cite-weighted quality.

5.53 Impact of Exogenous Shocks

The results presented so far have dealt with endogeniety issues firstly by including firm fixed effects (to deal with unobservables) and then by including a full set of 4-digit industry time trends (to capture trends in patenting). Both of these measures reduced the coefficient on defense sales across specifications. Table 5.5 reports further evidence based on the exogenous shocks variable outlined in section III. Panels (A) and (B) cover patents and the 'company-plus' measure of R&D respectively, with columns (1) and (4) repeating the analogous OLS models from earlier tables.

Column (2) reports the first stage reduced form equation for patents. This indicates that the product market shocks variable is a positive and highly significant predictor of firm procurement receipts. This carries through to the reduced form regression of product market shocks directly on patents reported in column (2). These results are consolidated in the IV estimate reported in column (3). Interestingly, the estimates here are higher than those for the equivalent OLS, within-groups specification (column(1)). A similar pattern of results holds for R&D in panel (B). Given the probable influence of the measurement error in defense sales (exacerbated by within groups estimation) the fact that the IV is estimate is higher than the OLS can be expected to some extent. Furthermore, the IV estimates in column (4) and (8) are still lower than the basic OLS estimates without fixed effects presented in Table 5.2. Finally, the IV estimates could conceivably be picking up the effects of spillovers from defense-induced increases in R&D and knowledge capital stocks that are external to the firm. I discuss this in more detail at the end of the next section.

5.54 Validity of IV Strategy

The validity of the overall IV result depends on the robustness of the identification strategy. One of the key assumptions of this strategy is that individual firms (or indeed groups of firms operating in the same product area) are unable to affect the level of spending chosen by the DoD. Specifically, it would be problematic if large, highly innovative firms were able to influence the pattern of the DoD's spending towards their own industries and product groups. Two pieces of evidence on DoD spending patterns are useful here. Firstly, Figure 4 shows how the share of spending for the Top 10 product groups has changed over time. As discussed, spending is very concentrated on this group of 10 products which unsurprisingly contain the major categories for the production of aircraft, ships, engines and electronics. The share of spending for the top 10 only begins to trend downwards after the end of the Reagan build-up in the late 1980s. Hence, stability in product shares up until this time was followed by a period where spending began to shift *away* from the technologically intensive Top 10 group. This is of course opposite to the expectation that firms in these industries could be influencing DoD spending policy in their own favor.

Table 5.6 provides a second piece of evidence that reinforces the message from Figure 4. This table is based on a collapsed 2-digit product panel of all DoD procurement spending. The objective here is to test whether spending changes are correlated with the level of market concentration in each product group. The independent variable used is a Herfindahl index, defined as:

$$H_{lt} = \sum_{i=1}^{N} h_{ilt}^2$$

Where h_{ilt} is the market share of a contractor selling goods to the DoD in product group *l*. These shares are calculated for the top 50 contractors in each 2-digit product group. Higher values of the index are associated with more concentrated product groups. The results indicate a negative associated between the Herfindahl index and changes in spending at the 2-digit level. This is consistent with the trend evident in Figure 4, namely that DoD spending has been trending strongly away from the concentrated, technologically-intensive product groups. This supports the assumption that firms in these industries are unable to affect the level of spending chosen by the DoD in their own favor[149].

[149] As discussed, another issue for this identification strategy is the potential role of spillovers that are correlated with firm-level spending shocks. This issue was outlined in detail in earlier sections and will implemented in future iterations of the paper.

Finally, another issue for the identification strategy is the potential role of spillovers. As outlined in Appendix A, the product market shocks variable could affecting firms indirectly via spillover channels as well as directly through increased firm procurement receipts. Put simply, any DoD demand shocks that impact a given firm *i* will affect other firms working in the same technology space as *i*. If these other firms then increase their R&D and patenting firm *i* could benefit following the standard spillovers argument. Hence the spillover provides an extra effect of defense procurement in addition to direct effect of firm receipts. The fact that the IV estimates are higher than the OLS bears this out. Hence, future iterations of this paper will test for these spillover effects following the approach outlined in Appendix A.1.

5.55 Historical Magnitudes

The results up until this point have focused on the firm-level relationship between technology and defense procurement sales. These results can be nested alongside an industry decomposition to understand the historical magnitude of the defense sector's contribution to innovation over different policy phases. Table 5.7 reports the results of a SIC3 industry decomposition of patenting and R&D intensity for each 4-year Presidential administration. As discussed, this decomposition follows the approach of Berman et al (1994) who studied the impact of trade, the defense build-up and biased technical change on the demand for skilled labor in US manufacturing. Note that the results in Table 5.7 are calculated using the full, collapsed COMPUSTAT database rather than just the sub-sample of firms who receive funding from defense contracts at some point. This is in line with the goal of estimating the defense sector's contribution with respect to the whole population of listed, corporate innovators.

Technological intensity grew strongly over the period being considered. The aggregate R&D share of sales grew from 0.014 in the 1977 to 0.024 in 1988 and 0.036 by 2000. Similarly, the number of patents per dollar of sales grew from approximately 0.036 per ten million dollars of industry sales in 1977 to 0.043 in 1988 and 0.075 in 2000. The first point to note from Table 5.7 is the overall breakdown of the within and between components behind this growth in R&D and patenting intensity. The first column indicates that between industry shifts account for around 35-45% which is comparable to other studies of this type[150]. Columns (2) and (3) then report the share of the between and within components that is accounted for by the defense sector. For

[150] For example, the Berman et al (1994) study of skilled labor shares estimates a total between component of 37% for the 1973-79 period and 30% for 1979-1987.

example, 5.5% of the total between share of 43.1% in 1981-1984 is due shifts in the defense sector. In turn, this yields the total defense contribution of 2.4% (column 4). In plain english, this says that the between industry shifts induced by the defense sector accounted for 2.4% of the total rise in patenting intensity in the 1981-1984 period.

The main qualitative message from Table 5.7 is the big role that the defense sector had in the first half of the Reagan build-up. The defense contribution peaked in the 1981-84 period, accounting for 11.4% of the change in the aggregate patenting intensity and 6.5% for R&D intensity. In the case of patenting, most of the effect (9%) is due to the within component. This is consistent with the findings of the firm-level models which showed that defense procurement stimulated within-firm patenting by more than would be expected from a conventional sales shock. The within industry component for R&D intensity (3.3%) is lower than the analogous figure for patenting and again this is consistent with the firm-level evidence of more limited R&D effects.

The figures for the later post-Cold War periods show the predicted impact of reductions in defense procurement on patenting and R&D intensity. Specifically, since we are dealing with changes over time, these figures show how the reduction in defense spending slowed down the growth in technological intensity. This effect is strongest for the period immediately after the Reagan build-up (1989-1992) where the between industry shift away from defense production had a negative contribution of around -2% across both measures of technology. Interestingly, these is a persistent, positive within- industry defense contribution for patenting in the 1989-1996 period. Finally, note that while the cuts to defense procurement slowed down the growth of technology this was offset at the aggregate level by surges of patenting and R&D investment in other parts of the economy[151].

5.6 CONCLUSION

In this paper I have examined the impact of defense procurement spending on innovation and other related outcomes in a long-run sample of US listed firms. The motivation is straightforward: defense procurement spending is one of the major, direct policy channels through which the government can affect firms. This has been the case since WWII and the sharp increase in defense spending that occurred after 9/11 makes it of continued relevance. The high-tech composition of procurement spending also makes it a significant *de facto* innovation policy,

[151] See Kortum and Lerner (1998) and the Congressional Budget Office (2005) for discussions of these trends for patenting and R&D investment respectively.

alongside more explicit policies such as R&D tax credits and government support for basic science.

The paper has put forward evidence on two main questions. First is the relative impact of defense procurement sales on firm patenting and R&D expenditure, using civilian sales as the benchmark. The results indicate that the elasticity between defense procurement and both measures of innovation is approximately 0.07. This elasticity is in line with the civilian benchmark for R&D but well above the same benchmark for patenting. Furthermore, the patents stimulated by defense procurement maintain their level of quality measured in terms of citations. Direct evidence on this question has not been available before. It has been often speculated that the narrow, mission-focus of defense spending may limit its potential impact on innovation. On the other hand, the high-tech composition of defense procurement sales was always likely to skew the innovative impact of this spending upwards. The results in this paper strongly support this second view of procurement as a source of demand for high-tech products.

However, it must be noted that the results currently compare average defense sales against average civilian sales, with no allowance for the composition of civilian spending. The availability of data on non-defense Federal contracts allows me to look at this issue in more detail[152]. Future iterations of this paper will therefore test the efficacy of defense sales against the sales of comparable goods purchased by other Federal departments. This will allow me test whether there is a distinctive effects of defense procurement after controlling for product composition.

The second major question discussed has been the magnitude of the impacts from defense procurement. The basic magnitudes calculated here indicate that defense procurement accounted for 6-11% of the growth in patenting and R&D during the early Reagan-build-up. Cutbacks in spending during the Bush and Clinton administrations then acted as moderating influence, slowing innovation by up to 2%. Finally, it must be noted that this finding on magnitudes should not be considered an endorsement of defense procurement as a primary innovation policy tool. The final ledger entry on procurement impacts needs to take account of cost-effectiveness and

[152] The National Archive holds also records for all non-defense Federal procurement contracts from 1979-97. After 1997 these contract records are kept as part of the Federal Procurement Data System (FPDS).

directly compare dollar-for-dollar impacts from procurement against policy tools such as R&D tax credits.

REFERENCE LIST

Acemoglu, D and Linn, J (2004) “Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry”. *Quarterly Journal of Economics*. 119(3):1049-1090.

Berman, E; Bound, Eli and Griliches, Z (1994) “Changes in the Demand for Skilled Labor Within US Manufacturing Industries: Evidence from the Annual Survey of Manufacturing”. *Quarterly Journal of Economics*. 109(2): 367-397.

Bloom, N.; Draca, M. and Van Reenen, J. (2011) “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, Diffusion and Productivity”. *CEP Discussion Paper*. No 1000, January.

Bloom, N.; Schankerman, M. and Van Reenen, J. (2011) “Identifying Technology Spillovers and Product Market Rivalry”. *Revised Mimeo,* LSE.

Bush, V (1946*) Science: The Endless Frontier*. Office of Scientific research and Development, Washington.

Congressional Budget Office (2005) *R&D and Productivity Growth*. Background Paper, June. CBO: Washington.

Cozzi, G and Impullitti, G (2010) “Government Spending Composition, Technical Change, and Wage Inequality”. *The Journal of the European Economic Association*. 8(6): 1325-1358.

David, P; Hall, B; and Toole, A(2000) “Is Public R&D a Complement or Substitute for Private R&D? A Review of the Econometric Evidence”. *Research Policy* 29: 497-529.

Finkelstein, A (2004) “Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry”. *Quarterly Journal of Economics*. 119(2):527-564.

Goolsbee, A (1998) "Does Government R&D Policy Mainly Benefit Scientists and Engineers?" *American Economic Review*. *88*(2): 298-302.

Griliches, Z (1979) “Issues in Assessing the Contribution of R&D to Productivity Growth”. *The Bell Journal of Economics* 10 (1): 92-116.

Grilches, Z (1992) “The Search for R&D Spillovers”. *The Scandinavian Journal of Economics*. 94(0): S29-47.

Office of Management and Budget (2008) *Historical Tables of the US Budget*. Government Printng Office, Washington.

Hines, J and Guthrie, S (2011) “U.S. Defense Contracts During the Tax Expenditure Battles of the 1980s”. *National Tax Journal*. 64(2): 731-52.

Hall, B. H.; Jaffe, A. B. and Trajtenberg, M. (2001) “The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," *NBER Working Papers* 8498, National Bureau of Economic Research, Inc.

Hausman, J; Hall, B and Griliches, Z (1984) “Econometric Models of Count Data with an Application to the Patents-R&D Relationship”. *Econometrica*. 52(4): 909-938.

Howe, J and McFetridge, D (1976) “The Determinants of R&D Expenditures”. *Canadian Journal of Economics*. 9(1): 57-71.

Jaffe, A (1986) “Technological Opportunity and Spillovers of R&D: Evidence from Firms Patents, Profits and Market Value”. *American Economic Review*. 76(5): 984-1001.

Kortum, S and Lerner, J (1998) “Stronger Protection or Technological Revolution: What is Behind the Recent Surge in Patenting?”. *Carnegie Rochester Conference Series on Public Policy*. 48: 247-304.

Le Large, C and Nefussi, B(2010) “The Dynamics of Firms’ Product Portfolio in Response to International Competition : An Empirical Assessment”. Mimeo, CREST-INSEE.

Lichtenberg, F. (1988) “The Private R&D Investment Response to Federal Design and Technical Competitions”. *American Economic Review* 78(3):550-559.

Lichtenberg, F. (1989) "IR&D Project Data and Theories of R&D Investment," *Journal of Economic Dynamics and Control*. 13: 271-82.

Lichtenberg, F. (1990b) “How Elastic is the Government’s Demand for Weapons?” *The Journal of Public Economics* 40(1):57-78.

Lichtenberg, F. (1990b) “U.S. Government Subsidies to Private Military R&D: The Defense Department's Independent R&D Policy”. *Defense Economics* (1): 149-58.

Lichtenberg, F. (1992) “A Perspective on Accounting for Defense Contracts”. *The Accounting Review*. 67: 741-752.

Lichtenberg, F and Siegel, D (1991) “The Impact of R&D Investment on Productivity: New Evidence Using Linked R&D-LRD Data”. *Economic Enquiry* 29(2): 203-228.

Lichtenberg, F. (1995) “The Economics of Defense R&D”. In eds Hartley, K and Sandler, T. *Handbook of Defense Economics, Volume 1*. Elsevier: New York.

Markusen, A.; Hall, P.; Campbell,S. and Deitrick,S. (1991) *The Rise of the Gunbelt: The Military Remapping of Industrial America*. Oxford University Press: Oxford.

Mowery, D (2010) “Military R&D and Innovation”. In Hall, B and Rosenberg, N (eds) *The Handbook of the Economics of Innovation*, Volume 2. Elsevier, London.

Mowery, D and Rosenberg, N (1991) *Technology and the Pursuit of Economic Growth*. Cambridge University Press: Cambridge, UK.

Nakamura, E and Steinsson, J (2011) “Fiscal Stimulus in a Monetary Union: Evidence from US Regions”. Mimeo, Columbia University.

National Science Foundation (2006) *National Patterns of R&D Resources (Annual Series)*. Division of Science Resources Statistics.

Newell, R.; Jaffe, A. and Stavins, R. (1999) “The Induced Innovation Hypothesis and Energy Saving Technological Change”. *Quarterly Journal of Economics* 114(3):907-940.

Peck, J and Scherer, D (1962) *The Weapons Acquisition Process: An Economic Analysis*. Harvard Business School, Boston, MA.

Popp, D. (2002) “Induced Innovation and Energy Prices”. *American Economic Review* 92(1): 160-180.

Rogerson, W. (1992) “Overhead Allocation and Incentives for Cost Minimization in Defense Procurement”. *The Accounting Review*. 67: 671-690.

Ruttan, V. (2006*) Is War Necessary for Economic Growyh? Military Procurement and Technology Development*. Oxford University Press: Oxford.

Schmookler, J. (1966) *Invention and Economic Growth.* Harvard University Press: Cambridge, MA.

Thomas, J and Tung,S (1992) “Incentives Under Cost-Reimbursement: Pension Costs for Defense Contractors”. *The Accounting Review*. 67: 671-690.

U.S. National Archives, Records of Prime Contracts Awarded by the Military Services and Agencies, various years (http://aad.archives.gov/aad/series-list.jsp?cat=GS29)

FIGURE 5.1: REAL DEFENSE SPENDING.

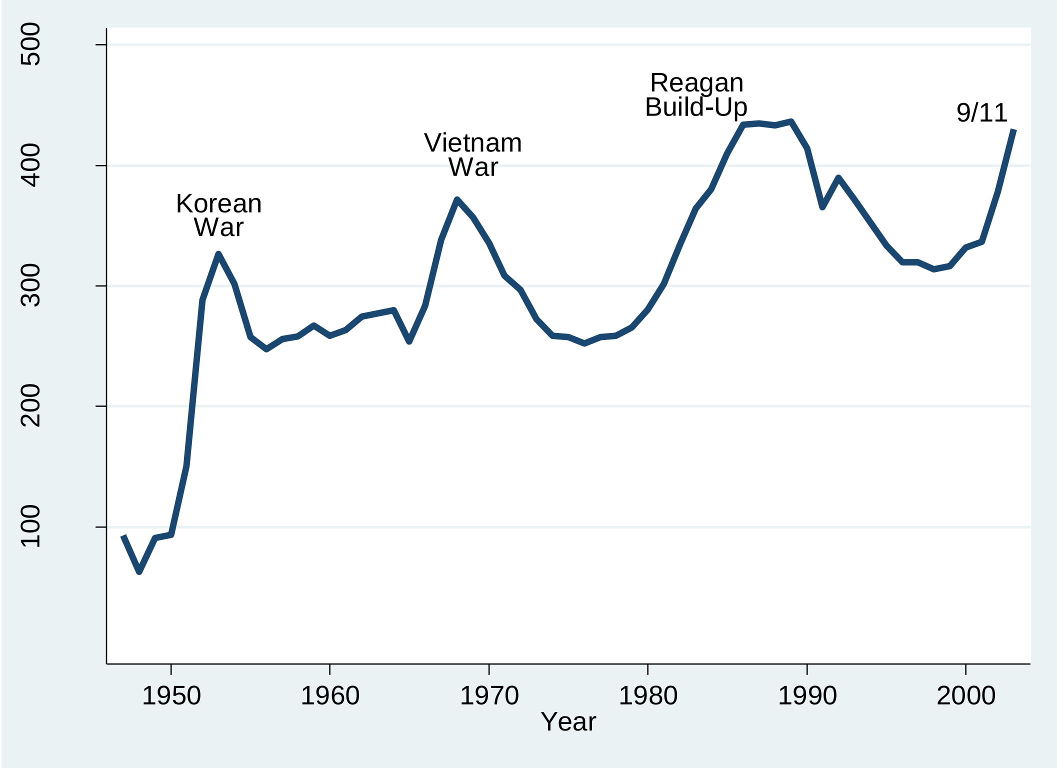

Notes: Source is Table 6 (Composition of Outlays) in *Historical Tables, Budget of the US Government, Fiscal Year 2005*. Deflated by the GNP deflator (base year 2003).

FIGURE 5.2A: TOTAL DEFENSE PROCUREMENT RECEIPTS, FY1966-FY2003.

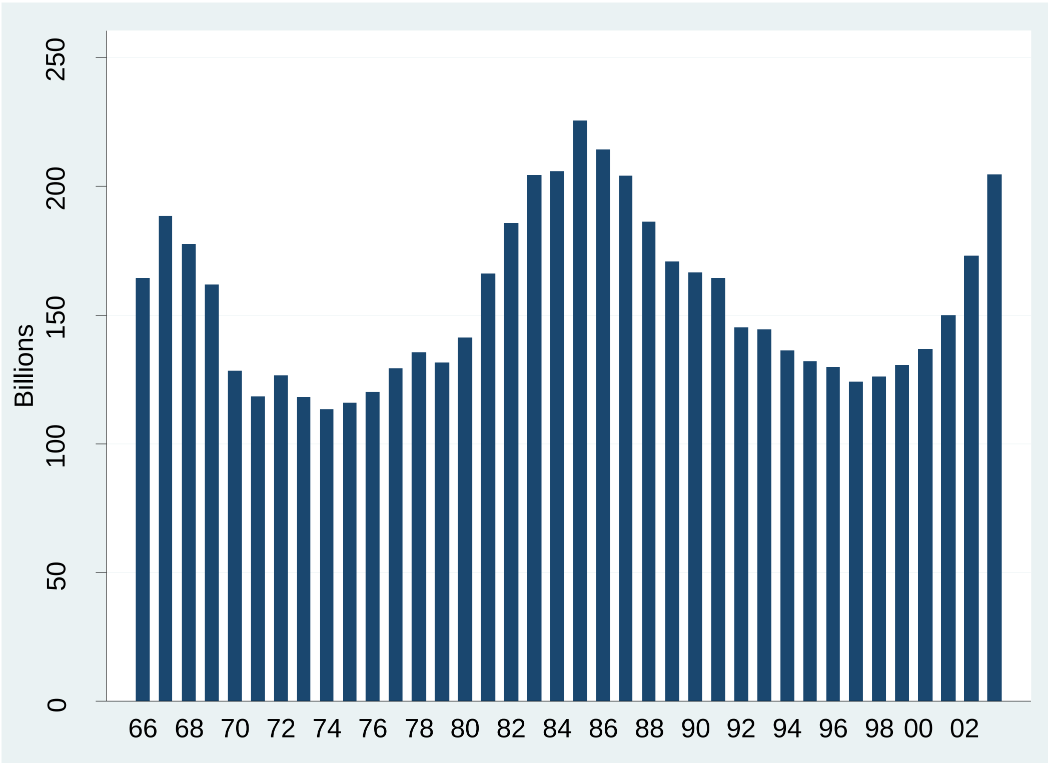

Source: National Archives and Records Administration (NARA) Historical Files on military procurement. Deflated by the GNP deflator (base year 2003).

FIGURE 5.2b: PROCURMENT SPENDING SHARES, TOP 10 2-DIGIT PRODUCTS.

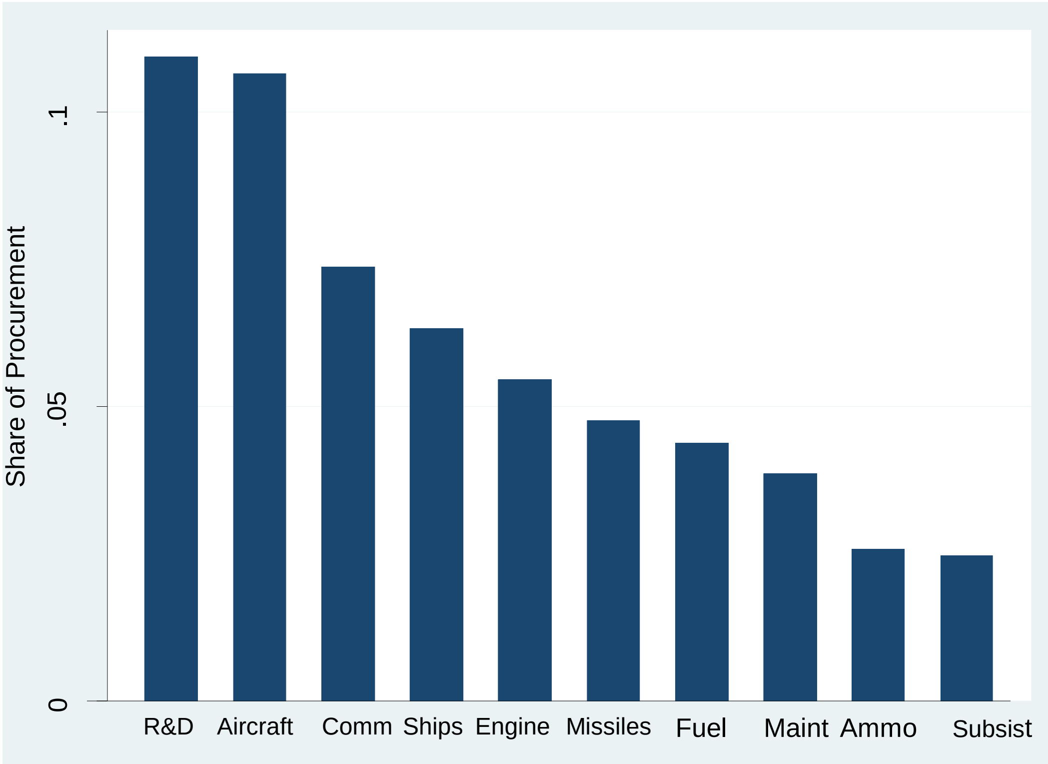

Notes: These top 10 products are: R&D (for Equipment); Aircraft & Airframe Structural Components; Communications, Detection and Coherent Radiation equipment; Ships and Small Craft; Engines, Turbines and Components; Guided Missiles; Fuel and Lubricants; Ammunition & Explosives; Maintenance and Repair; and Subsistence. Note that the R&D (for Equipment) category includes R&D on Aircraft, Missiles, Ships, Tanks, Weapons and Electronics. Source is National Archives and Records Administration (NARA) Historical Files on military procurement.

FIGURE 5.3: DETERMINANTS OF PRIVATE R&D INVESTMENT.

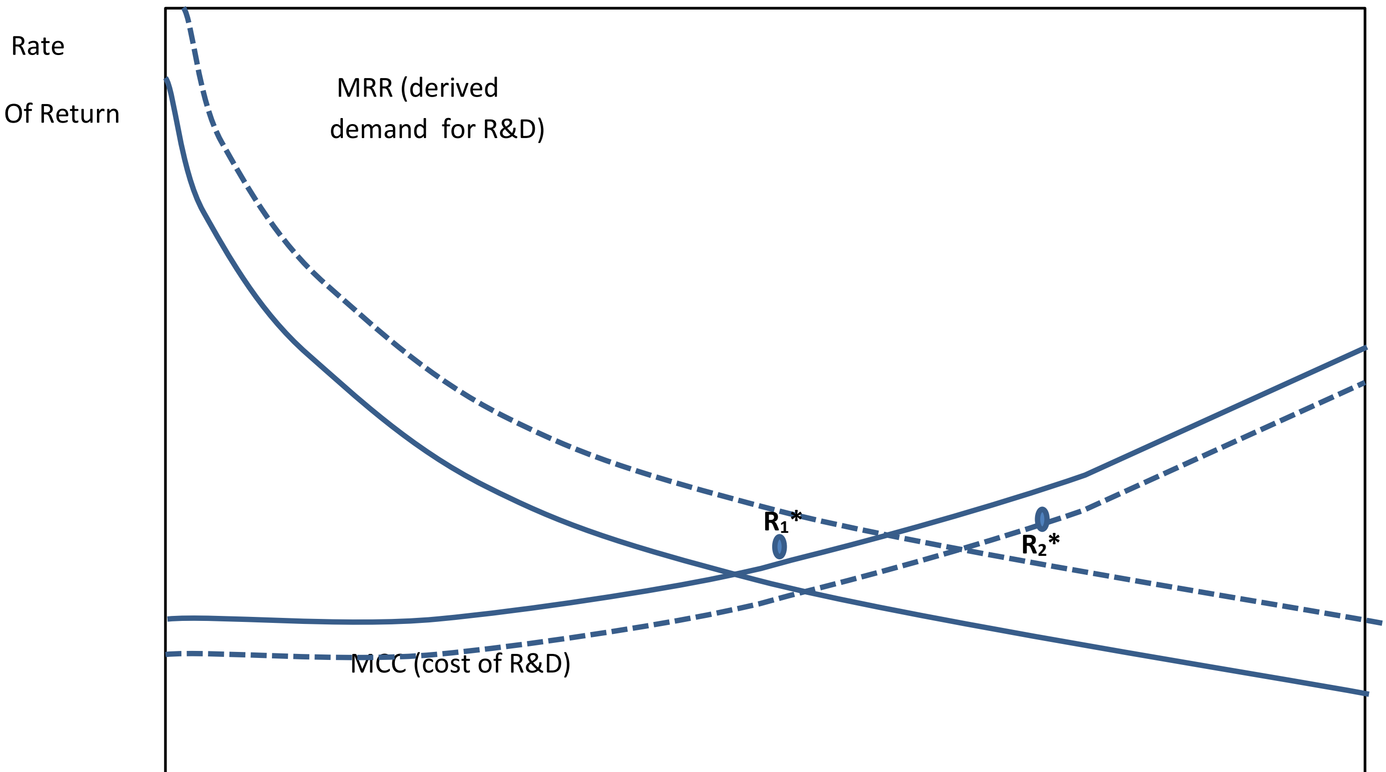

Notes: This figure illustrates the discussion in Section II (Background). Defense contractors receive a subsidy as part of the Independent Research and Development (IR&D) program to assist them in participating technical bids for DoD design competitions. This subsidy lowers the cost of capital shifts the MCC schedule down. Procurement contracts (often occurring as non-competitive follow-ons from design competitions) shift the MRR schedule outwards and are quantitatively more valuable than the IR&D subsidy.

FIGURE 5.4: SHARE OF TOP 10 PRODUCTS IN DEFENSE PROCURMENT, FY1996-FY2003

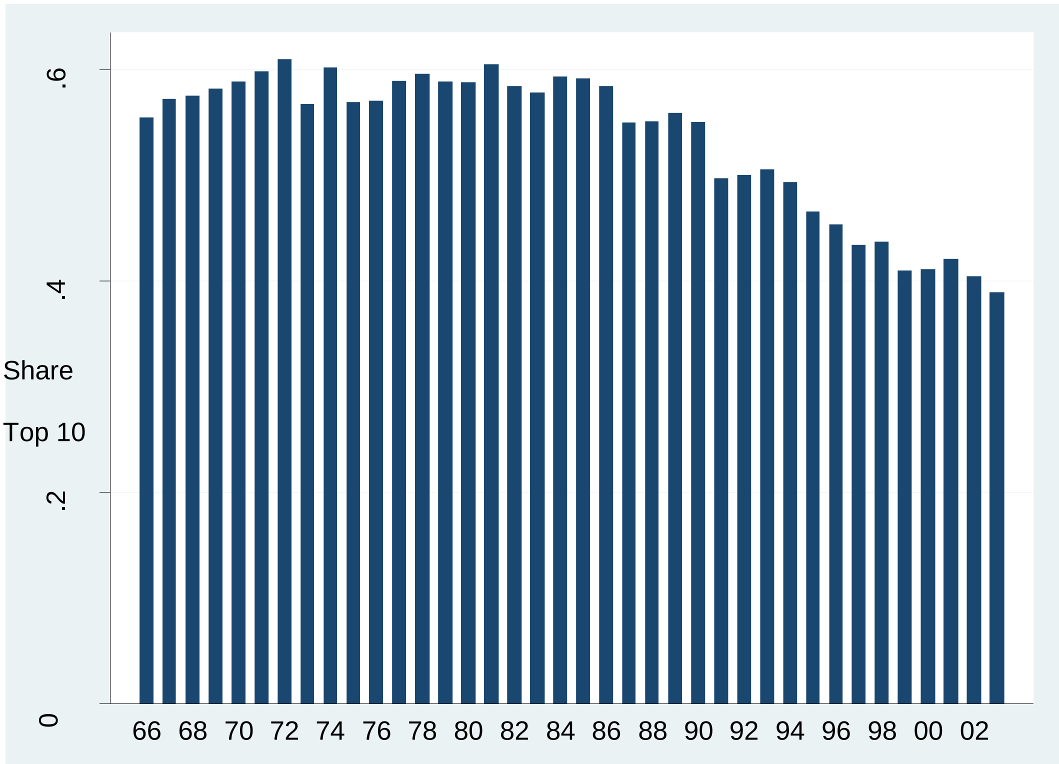

Notes: This figure shows the share of the 'Top 10' 2-digit products by dollar value in total procurement spending. These top 10 products are: Aircraft & Airframe Structural Components; Ships and Small Craft; Engines, Turbines and Components; Communications, Detection and Coherent Radiation; Aircraft R&D; Ammunition & Explosives; Guided Missiles; Subsistence; Fuel and Lubricants; and Maintenance and Repair. Source is National Archives and Records Administration (NARA) Historical Files on military procurement.

TABLE 5.1: DESCRIPTIVE STATISTICS, MATCHED SAMPLE, 1966-2003 (MANUFACTURING)

	(1) No Defense Procurement	(2) Receive Defense Procurement (Some Years)	(3) Receive Defense Procurement (All Years)
Patent Count	1.9	10.4	24.0
	(21.4)	(80.0)	(87.6)
Citation Count	11.4	82.9	216.5
	(121.8)	(656.2)	(828.6)
‘Military’ Patent Count	0.7	4.0	7.2
	(11.5)	(33.0)	(24.8)
Employment (in 1000s)	2.9	5.1	11.3
	(8.2)	(12.2)	(18.1)
Employment (median)	0.51	1.1	3.1
Sales	648.4	1073.3	2272.0
	(2133.0)	(2953.8)	(4115.8)
Sales (median)	93.9	181.2	439.4
Company-sponsored R&D	18.2	39.1	73.1
	(79.3)	(111.3)	(170.9)
Defense R&D	na	2.6	12.0
		(28.8)	(55.3)
(R&D/Sales)	0.045	0.040	0.040
	(0.123)	(0.095)	(0.060)
(Defense Sales / Sales)	na	0.014	0.049
		(0.071)	(0.116)
Defense Sales (DS)	na	7.0	52.4
		(106.5)	(238.3)
SIC35 Share	0.108	0.159	0.188
SIC36 Share	0.128	0.156	0.217
SIC37 Share	0.043	0.054	0.080
SIC38 Share	0.088	0.125	0.174
Number of Firms	5,976	2,207	664
Number of Observations	56,394	36,270	19,579

Column (1) reports statistics for firms who never receive any defense procurement contracts, Column (2) reports for firms who receive contracts in some but not all years. Column(3) reports for the sample of firms who receive contracts in all observed years. Sales, R&D, Defense Sales in million dollar units.

TABLE 5.2: PATENTING, R&D AND DEFENSE PROCUREMENT, 1976-2003.

(A) Patenting	(1) ln(1+PAT)	(2) ln(1+PAT)	(3) ln(1+PAT)	(4) (PAT)
ln(Defense)$_{t-1}$	0.320*** (0.029)	0.079*** (0.016)	0.071*** (0.014)	0.080*** (0.008)
Firm Fixed Effects	No	Yes	Yes	Yes
SIC4*Year Trends	No	No	Yes	Yes
Number of Firms	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116

(B) R&D	(1) ln(CR&D)	(2) ln(CR&D)	(3) ln(CR&D)	(4) ln(GR&D)
ln(Defense)$_{t-1}$	0.409*** (0.029)	0.071*** (0.016)	0.065*** (0.012)	0.095** (0.014)
Firm Fixed Effects	No	Yes	Yes	Yes
SIC4*Year Trends	No	No	Yes	Yes
Number of Firms	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116

Notes: Standard errors clustered by firm in parentheses. Column (1) includes SIC4 fixed effects. The dependent variable ln(1+PAT) is the log of 1 plus the patent count (patents applied for by firm *i* in year *t*). The variable ln(CR&D) is the log of company-sponsored R&D expenditures by firm *i* in year *t*, as reported in COMPUSTAT. The variable GR&D represents the sum of company-sponsored R&D reported in COMPUSTAT plus the sum of Department of Defense procurement-funded R&D reported in the NARA files for firm *i* in year *t*. ln(Defense Sales)$_{t-1}$ is the log of the total value of procurement contracts received by the firm *i* in year *t-1*. Column (4) is uses the count of patents as the dependent variable with the model estimated using a negative binomial conditional MLE.

TABLE 5.3: BENCHMARKING AGAINST CIVILIAN SALES, 1976-2003.

(A) Patenting	(1)	(2)	(3)	(4)
	ln(1+PAT)	ln(1+PAT)	ln(1+PAT)	(PAT)
$\ln(\text{Defense})_{t-1}$	0.069***	0.060***	0.059***	0.041***
	(0.019)	(0.015)	(0.012)	(0.007)
$\ln(\text{Civilian})_{t-1}$	0.638***	0.446***	0.396***	0.303***
	(0.026)	(0.051)	(0.044)	(0.014)
Implied $\square_1$	0.038	0.027	0.024	0.018
Test $\square_1$ (p-value)	0.10	0.028	0.003	0.606
Firm Fixed Effects	No	Yes	Yes	Yes
SIC4*Year Trends	No	No	Yes	No
Number of Firms	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116

(B) R&D	(1)	(2)	(3)	(4)
	ln(CR&D)	ln(CR&D)	ln(CR&D)	ln(GR&D)
$\ln(\text{Defense})_{t-1}$	0.040***	0.039***	0.045***	0.077***
	(0.013)	(0.011)	(0.009)	(0.012)
$\ln(\text{Civilian})_{t-1}$	0.941***	0.751***	0.646***	0.600***
	(0.018)	(0.043)	(0.039)	(0.040)
Implied $\square_1$	0.056	0.045	0.038	0.036
Test $\square_1$ (p-value)	0.215	0.59	0.499	0.005
Firm Fixed Effects	No	Yes	Yes	Yes
SIC4*Year Trends	No	No	Yes	Yes
Number of Firms	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116

Notes: Standard errors clustered by firm in parentheses. Column (1) includes SIC4 fixed effects. The dependent variable ln(1+PAT) is the log of 1 plus the patent count (patents applied for by firm *i* in year *t*). The variable (PAT) is the count of patents applied for by firm *i* in year *t*. The variable ln(CR&D) is the log of company sponsored R&D expenditures by firm *i* in year *t*, as reported in COMPUSTAT. The variable GR&D represents the sum of company-sponsored R&D reported in COMPUSTAT plus the sum of Department of Defense procurement-funded R&D reported in the NARA files for firm *i* in year *t*. $\ln(\text{Defense Sales})_{t-1}$ is the log of the total value of procurement contracts received by the firm *i* in year *t-1*. ln(Civilian Sales) is the log of 'Civilian' non-defense sales calculated as total reported accounting sales (from COMPUSTAT) minus the total value of procurement contracts given in the NARA files (ie: Defense Sales). The implied $\square_1$ is calculated as $\square_2 * (D / C)$, the coefficient on Civilian sales multiplied by the ratio of defense to civilian sales. The (D/C) ratio for this sample is 0.060.

TABLE 5.4: TYPES OF PATENTS AND CITATIONS, 1976-2003.

	(A) Patents			(B) Citations		
	(1)	(2)	(3)	(4)	(5)	(6)
	ln(1+PAT)	1n(1+MPAT)	1n(1+CPAT)	ln(1+CITE)	ln(1+MCITE)	ln(1+CCITE)
ln(Defense)$_{t-1}$	0.059***	0.048***	0.050***	0.067***	0.085***	0.063***
	(0.012)	(0.010)	(0.012)	(0.021)	(0.022)	(0.020)
ln(Civilian)$_{t-1}$	0.396***	0.268***	0.352***	0.532***	0.435***	0.496***
	(0.044)	(0.036)	(0.042)	(0.065)	(0.061)	(0.065)
Implied $\square_1$	0.024	0.016	0.021	0.032	0.026	0.030
Test $\square_1$ (p-value)	0.004	0.003	0.011	0.080	0.007	0.097
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
SIC4*Year Trends	Yes	Yes	Yes	Yes	Yes	Yes
Number of Firms	664	664	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116	7,116	7,116

Notes: Standard errors clustered by firm in parentheses. The dependent variable ln(1+PAT) is the log of 1 plus the patent count (patents applied for by firm *i* in year *t*). The variable ln(1+MPAT) is the log of 1 plus the count of patents belonging to military patent classes following the discussion in section III ("Patents Data – Defining Military Patents"). The variable ln(1+CPAT) represents the count of patents in all the remaining non-military or 'civilian' patent classes. The variable ln(1+CITE) is the log of 1 plus the count of citations (ie: forward citations for the patents applied for by firm *i* in year *t*). The variables ln(1+DCITE) and ln(1+CCITE) are then the defense and civilian analogues of the total citations variable.

TABLE 5.5: IMPACT OF DEFENSE PRODUCT MARKET SHOCKS – REDUCED FORM AND IV ESTIMATES, 1976-2003.

	(A) Patenting				(B) R&D			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	Reduced Form	First Stage	IV	OLS	Reduced Form	First Stage	IV
	ln(1+PAT)	ln(1+PAT)	$\ln(\text{Defense})_{t-1}$	ln(1+PAT)	ln(R&D)	ln(1+PAT)	$\ln(\text{Defense})_{t-1}$	ln(R&D)
$\ln(\text{Defense})_{t-1}$	0.071***			0.173**	0.095***			0.265**
	(0.014)			(0.077)	(0.014)			(0.077)
$\ln(\sum \Phi_{ij,(t-10)} D_t^L)$		0.022**	0.130***			0.034**	0.130***	
		(0.011)	(0.020)			(0.010)	(0.020)	
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SIC4*Year Trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of Firms	664	664	664	664	664	664	664	664
Number of Observations	7,116	7,116	7,116	7,116	7,116	7,116	7,116	7,116

Notes: Standard errors clustered by firm in parentheses. Estimated for the period 1976-2003 (Fiscal Years) where a 10-year lag of the firm-specific defense shift-share variable is defined. The dependent variable ln(1+PAT) is the log of 1 plus the patent count (patents applied for by firm *i* in year *t*) . $\ln(\text{Defense Sales})_{t-1}$ is the log of the total value of procurement contracts received by the firm *i* in year *t-1* . The variable $\ln(\sum \Phi_{ij,(t-10)} D_t^L)$ is the log of the firm-specific Department of Defense (DoD) product market. This is defined as the firm's share of own sales to the DoD in each 2-digit Federal Supply Code (FSC) 10 years ago multiplied by the current, period (t-1) value of DoD spending in each 2-digit category.

TABLE 5.6: CHANGES IN DOD SPENDING AND PRODUCT MARKET COMPETITION, 2-DIGIT PRODUCT PANEL,1976-2003.

	(1)	(2)	(3)
Dependent Variable	$\Delta \ln(D^L)_{lt}$	$\Delta \ln(D^L)_{lt}$	$\Delta \ln(D^L)_{lt}$
$(\text{Herfindahl})_{t-2}$	-0.087		
	(0.092)		
$(\text{Herfindahl})_{t-5}$		-0.061	
		(0.059)	
$(\text{Herfindahl})_{t-10}$			-0.024
			(0.066)
1-Digit Trends	Yes	Yes	Yes
Number of Product Groups	154	154	154
Number of Observations	3353	3353	3353

Notes: Standard errors clustered by 2-digit product category in parentheses. The data is a 2-digit product level panel of NARA procurement data. The dependent variable is the log 1-year change in the total amount of procurement spending at the 2-digit level. The Herfindahl index is calculated as the sum of the squared market shares for the top 50 contractors in each 2-digit product group. There are 34 1-digit groups with a trend term included for each group.

TABLE 5.7: INDUSTRY SECTOR DECOMPOSITION, 1977-2000.

Panel (A) Patents/Sales

Period	(1) Total Share Of Between	(2) Defense Share Of Between	(3) Defense Share Of Within	(4) Defense Contribution Between	(5) Defense Contribution Within	(6) Defense Contribution Total
Pre Build-Up						
1977-1980	0.353	0.041	0.029	0.014	0.019	0.033
Reagan Build-Up						
1981-1984	0.431	0.055	0.159	0.024	0.090	0.114
1985-1988	0.381	0.026	-0.001	0.010	-0.001	0.009
Post Cold War						
1989-1992	0.378	-0.044	0.022	-0.017	0.014	-0.003
1993-1996	0.370	-0.019	0.014	-0.010	0.009	0.002
1997-2000	0.344	-0.007	0.004	-0.002	0.003	0.001

Panel (B) R&D/Sales

Period	(1) Total Share Of Between	(2) Defense Share Of Between	(3) Defense Share Of Within	(4) Defense Contribution Between	(5) Defense Contribution Within	(6) Defense Contribution Total
Pre Build-Up						
1977-1980	0.43	0.031	0.030	0.013	0.017	0.030
Reagan Build-Up						
1981-1984	0.348	0.093	0.050	0.032	0.033	0.065
1985-1988	0.391	0.028	0.032	0.011	0.019	0.030
Post Cold War						
1989-1992	0.517	-0.040	-0.021	-0.021	-0.010	-0.031
1993-1996	0.456	-0.020	-0.007	-0.009	-0.004	-0.013
1997-2000	0.450	-0.005	0.006	-0.002	0.003	0.001

Notes: This table reports the result of a SIC3 level decomposition of the growth in aggregate patenting and R&D expenditure. The full COMPUSTAT database is used with sales, patents and R&D collapsed to the SIC3 level, as defined by the main SIC code for each listed firm. The technology measures (patents and R&D) are normalized by industry sales to construct measures of technological intensity. Column (1) reports what share of the total change in intensity is due to between effects with column (2) then reporting what share of this between component is accounted for by the defense sub-sector. Column (4) then calculates the total contribution of the defense sector to the aggregate change in technological intensity. For example, in 1981-1984 2.4% of total change in patenting intensity is due to between industry shifts associated with the defense sub-sector. See section IV ('Calculating Magnitudes') for a full description of the decomposition.

APPENDIX 5.A.1

Spillovers in the Defense Product Market

The above shift-share approach to identifying exogenous shocks to procurement spending is complicated by the potential role of knowledge spillovers at the levels of industry, product or technological class. Empirically, spillovers are best described as a type of 'outside' capital that benefits the firm (Griliches 1992). The typical approach is to add up the stock of knowledge capital (measured most commonly by R&D) according to an external criteria. This basic scheme can be formulated as:

$$SPILL = \sum_{-i} v_{(i,-i)} G_{-i} \tag{9}$$

where $v_{(i,-i)}$ represents the 'distance' between firm *i* and all other firms *–i* and G_{-i} is the value of the knowledge capital for firm *–i*. In this paper's setting distance is a weight on the *–i* firms that can be measured in terms of industry (*j*), patent technological class (*k*) or DoD product code (*l*). If a *–i* firm is in the same industry, technological or product market space as firm *i* then it has a non-zero weight. Following the approach of Jaffe(1986) and latterly Bloom, Schankerman and Van Reenen (2011) we can define a measure of closeness in the defense product code space as:

$$PROD_{i,-i} = \frac{\left(D_i D_{-i}'\right)}{\left(D_i D_i'\right)\left(D_{-i} D_{-i}'\right)} \tag{10}$$

where $D_i = \left(d_{i1}, d_{i2}, ...d_{i120}\right)$ is the vector of firm *i*'s share of product *l* in its total defense sales across the 155 2-digit product codes. The term in (8) is therefore the uncentred correlation between pairings of firm *i* and all other -*i*. The spillover pool is then constructed by applying these weights and summing over the relevant firms:

$$SPILLPROD_{it} = \sum_{-i} PROD_{(i,-i)} G_{(-i)t} \tag{11}$$

The main issue here is that both the spillover weights and the share vector for our exogenous shocks term in (6) are defined in the defense product space. Consider the simple case of a firm *i* that specializes in one defense product *l* over time. Then any DoD spending shock in category *l* affects both the direct procurement receipts of firm *i* as well as the receipts of all firms also selling product *l*. If all of these firms invest in R&D as a consequence of the shock this will increase the accumulated R&D stock in product space *l*. It is then possible that this additional induced R&D at the product space level will go on to affect firm *i* through a knowledge spillover mechanism. In principle, this could lead to upward bias on the proposed IV estimate since it could pick up the indirect spillover effects along with the direct effects of procurement that are the main focus of the paper. Practically, this problem can be mitigated by conditioning on extra spillover terms when estimating the main models.

APPENDIX 5.A.2

EXAMPLE OF FULL TEXT PATENT SUBJECT TO GOVERNMENT SUPPORT

Method and apparatus for a reprogrammable program missile memory

AbstractA method and apparatus are disclosed for a reprogrammable program missile memory module 26 which is placed within a missile 14 in substantially the same manner as the currently used programmable read only memory. The reprogrammable program memory module 26 provides for remote writing of tactical program data thereto while allowing the missile 14 to remain in a substantially operational state.

Inventors: **Siering; Erik R.** (Woodland Hills, CA)

Assignee: **Hughes Aircraft Company** (Los Angeles, CA)

Appl. No.: **07/437,044**

Filed: **November 15, 1989**

Current U.S. Class:	**244/3.15** ; 711/103
Current International Class:	F41G 7/00 (20060101); G06F 001/00 ()
Field of Search:	364/900 244/3.15,3.11

References

U.S. Patent Documents

4037202	July 1977	Terzian
4660170	April 1987	Hui et al.
4935881	June 1990	Lowenson et al.

Primary Examiner: Hellner; Mark
Attorney, Agent or Firm: Brown; C. D. Heald; R. M. Denson-Low; W. K.

Government Interests

This invention was made with Government support under Contract No. N00019-85-G-0171 awarded by the Department of the ***Navy***. The Government has certain rights in this invention.