

Speculative Contracts*

Kfir Eliaz[†] and Ran Spiegler[‡]

First Version: November 2005

This Version: August 2006

Abstract

In many principal-agent environments, the two parties hold different prior beliefs regarding the agent's future preferences. These differences may be due to inherent biases such as over-optimism or over-pessimism. We analyze the principal's optimal contract design under the assumption that the agent's prior is private information. In order to screen the agent's prior, the principal devises a menu of contingent contracts, some of which are "speculative" as they involve betting on the agent's future action. We characterize the optimal menu and show that the characterization enables us to interpret real-life contract design in a variety of economic contexts.

1 Introduction

When designing the terms of a bilateral contract, parties need to take into account differences in opinion regarding the likelihood of future events. While standard models assume that these differences are a result of informational asymmetries, they could also be due to heterogeneity in prior beliefs. An entrepreneur seeking to finance a new project may be more optimistic than an investor about its prospects; a sales person may be more optimistic about his salesmanship than a prospective employer; an advertising agency and a client may disagree over which type of campaign would be most successful; a project manager and a contractor may disagree over which variety of a product will be desired by consumers; etc.¹

In principal-agent relations, the principal may exploit such differences in prior beliefs by writing a contingent contract, which is essentially a bet on the agent's future action. The reason parties may agree to bet is that each of them is willing to make a concession

*Financial support from the US-Israel Binational Science Foundation, Grant No. 2002298 is gratefully acknowledged.

[†]Dept. of Economics, NYU. 269 Mercer St., New York, NY 10003, E-mail: kfir.eliaz@nyu.edu, URL: <http://homepages.nyu.edu/~ke7>

[‡]School of Economics, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: rani@post.tau.ac.il. URL: <http://www.tau.ac.il/~rani>

¹For other examples, see Bazerman and Gillespie (1999).

in the contingency that he deems less likely. The traditional view of contingent contracts in economics is either as a tool for screening agents according to their preference types, or as a tool for overcoming moral hazard. However, when the two parties have different prior beliefs, contingent contracts serve as bets. We pose the following question: how should such a contract be designed when the agent's prior belief is private information?

We study a two-period principal-agent contracting model. The principal enables the agent to choose in period 2 from some set of actions, conditional on signing a contract in period 1. If the agent refuses to sign a contract, he chooses some outside option. A contract is a function that specifies a monetary transfer between the principal and the agent for every second-period action. Other than that, we place no restriction on the space of contracts. The agent has quasi-linear vNM utility over action-transfer pairs. However, his utility from actions is either u or v , depending on a state of nature, which is revealed to the agent alone in period 2. The principal assigns probability p to state u , while the agent assigns probability θ to this state. The principal does not observe θ , and believes that it is drawn from some distribution on $[0, 1]$. Hence, θ plays the role of an agent "type". The principal's problem is to design a *menu of contracts* that maximizes his expected profit.

The following situation illustrates the model. A hungry customer enters a restaurant, unsure of how much he will have to eat in order to curb his hunger. In one state (the "hungry" state), he will have to eat both a main course and a dessert in order to feel satiated. In the other state (the "satiated" state), a main course will suffice.² The restaurant manager believes that the states are equally likely. While the cost of preparing each course is \$5, the customer is willing to pay up to \$20 for a meal that curbs his hunger, and only \$10 for a course that leaves him hungry. Stated differently, the customer's willingness to pay for a main course is \$10 in the hungry state and \$20 in the satiated state.

If the customer was known to share the manager's prior, the latter could offer him an "all you can eat" buffet for a fixed price of \$20. This contract extracts the entire consumer surplus in each state, and guarantees the manager an expected profit of \$12.5.

Suppose next that the manager knows that the customer *underestimates* his hunger - specifically, that he assigns probability $\frac{3}{4}$ to the satiated state. The manager could still offer the above "all you can eat" buffet, which would be accepted by the customer and generate the same expected profit. However, another possibility would be to offer the customer an "a la carte" deal, in which the prices of a main course and a dessert are \$17.50 and \$10, respectively. If the customer accepts this menu, then in the hungry state he will order both a main course and a dessert, while in the satiated state he will order only a main course. The customer's expected payoff from this menu - calculated according to *his own* prior belief - is zero, and so he would accept it. As for the manager, this contract is superior to the "all you can eat" buffet because it generates an expected profit - calculated according to the

²Assume that the consumer can have dessert only after eating a main course ("you can't have your pudding if you don't eat your meat...").

manager's prior belief - of \$15.

Now suppose that the manager knows that the customer *overestimates* his hunger - specifically, that he assigns probability $\frac{3}{4}$ to the hungry state. In contrast to the previous case, an “a la carte” menu cannot generate a higher expected profit for the restaurant manager than the above “all you can eat” buffet. To see why, note that regardless of the customer’s prior belief, he is never willing to pay more than \$10 for a dessert. And since the customer type in question deems the hungry state more likely than the satiated state, he is willing to pay less than \$15 for a main course. Therefore, given the manager’s beliefs, the expected profit generated by any “a la carte” menu falls below \$12.5.

The “a la carte” deal offered to the customer who underestimates his hunger is a “speculative contract” - namely, a bet on whether customer will order dessert. A customer who shares the manager’s beliefs would never accept it. Note that if the manager does *not* know which of the three customer types he is facing, he could simply ask the customer to choose between the “all you can eat” and “a la carte” deals. The customer who underestimates his hunger would weakly prefer the speculative “a la carte deal”, while the other two types would strictly prefer the non-speculative, “all-you-can-eat” buffet.

This example has several noteworthy features. First, the optimal menu of contingent contracts contains a risky scheme as well as a scheme that guarantees the agent his reservation value in each state. Second, the customer’s behavior is independent of his choice of contract: he eats both a main course and a dessert in the hungry state, while avoiding dessert in the satiated state. Thus, the multiplicity of contracts has nothing to do with designing second-period incentives. Rather, its objective is to screen the customer’s prior belief. Third, the customer types who assign higher probability to the hungry state than the manager exert no informational externality on the customer type who underestimates his hunger. This allows the restaurant manager to design a speculative contract for the latter type, without worrying that this contract could be chosen by the other types. Finally, the manager chooses to speculate only with the customer type who is “optimistic” about the amount of food that he has to eat in order to curb his hunger.

How general are these effects? In Section 4, we characterize the optimal menu of contracts for a large class of utility functions u and v , having the methodologically useful property (which is explained in Section 3) that when prior beliefs are common (yet privately known by the agent), the optimal menu displays *no discrimination* between agent types. The only restriction we impose on the probability distribution from which the agent type is drawn (according to the principal’s second-order belief) is that it is continuous over $[0, 1]$. Under these assumptions, the key features of the optimal menu are as follows.

The optimal menu is the union of two distinct menus. The principal divides the problem of designing the optimal menu into two separate problems: the optimal menu for those agents whose prior on u is higher than his, and the optimal menu for those agents whose prior on v is higher than his. The two menus can be designed in such a way that the only incentive

constraints that the principal needs to worry about are those within each category: agents in one category have no incentive to pretend to belong to the other category (e.g., an “optimist” has no incentive to pretend to be a “pessimist”).

Speculative contracts. The contracts offered to agents whose beliefs are significantly different from the principal’s have a speculative component. Had these agents shared the principal’s prior, they would not have accepted the contracts offered to them.

Exclusion of bets. The optimal menu does not exclude agent types in the usual sense of failing to transact with them. Nevertheless, there *is* exclusion, in the sense that the principal does not bet with agents whose beliefs are close to his. The contract offered to these agents is the optimal contract when priors are common, and it involves no speculation.

Betting with optimists. When u weakly lies above v , the agent’s type can be interpreted as his “degree of optimism”. In this case, the optimal menu assigns a speculative contract only to agents whose prior on u is (significantly) *higher* than the principal’s. In other words, the principal never speculates with “pessimists”.

Ex-post efficiency. The contract chosen by each agent type induces him to choose the efficient action in the state he finds less likely than the principal. In Section 4 we identify a sufficient condition on u and v that guarantees ex-post efficiency in both states. In these cases, all agent types end up displaying the same state-dependent behavior, yet they choose different contracts and therefore make different payments. This effect is manifestly at odds with a model based on common priors.

Section 5 provides examples that demonstrate the relevance of our framework to real-life economic settings. First, we analyze an interaction between a manager and a prospective sales person, who are uncertain of the latter’s cost of generating sales. We demonstrate how our model explains the coexistence of fixed-wage and performance-based compensation schemes as a consequence of the manager’s attempt to screen the sales person’s degree of optimism. Next, we study in detail the above “restaurant” example and show how our model can give rise to menus that contain unlimited consumption contracts side by side with variable-rate contracts. We use this example to interpret deals offered by mobile phone companies and DVD rental stores.

Finally, in subsection 5.3 we examine another buyer-seller example, in which the two parties have different priors regarding the buyer’s future ideal variety of the seller’s product. We use this example to illustrate how speculative contracts may have an adverse effect on ex-post efficiency, such that the larger the gap between the parties’ prior beliefs, the greater the efficiency loss. In such environments, a central planner who cares about ex-post efficiency would prefer that the principal did not know the agent’s prior belief. In this sense, the asymmetric-information environment may be socially desirable relative to the complete-information environment, because it mitigates distorting effects of speculative contracting.

Related literature

Our paper is related to the contracting literature that studies non-linear pricing with imperfectly informed consumers (most notably, Baron and Besanko (1984), Armstrong (1996) and Courty and Li (2000)), a.k.a. “sequential screening”. Both this literature and our work study the problem of a monopolist who offers consumers a menu of contingent pricing schedules. In both frameworks, the monopolist’s objective is to screen agents according to their unobservable prior beliefs over their future tastes. However, in contrast to our model, the sequential screening literature assumes common priors. More specifically, it is common knowledge that the agent is better informed about his future tastes than the principal. As we show in Section 3, we focus on an environment for which this distinction is crucial: non-common priors are *necessary* to generate price discrimination in our environment.

We are aware of at least one precedent in the literature for a principal-agent model with non-common priors. Moscarini and Fang (2005) study the implication of non-common priors on the design of wage contracts. Unlike us, they analyze contract design with an informed principal, and focus on the signaling aspect of contracts. More specifically, they ask how a principal should design a wage contract when he holds the correct prior about his workers’ ability, whereas their priors are biased upwards. A key assumption in the paper is that an optimistic belief has a positive effect on a worker’s productivity. The principal, therefore, faces the following trade-off. On the one hand, he would like to provide the appropriate monetary incentives for his workers to exert effort. On the other hand, he is concerned that workers may infer their true ability from the contract he offers (what the authors call “morale hazard”).

While this paper is concerned with monopolistic screening of an agent’s prior belief, there have been a few works on competitive screening of priors. Landier and Thesmar (2004) examine debt contracts that are signed between investors and entrepreneurs who differ in their degree of optimism. Assuming a competitive environment, in which investors earn zero profits, the authors construct a separating equilibrium in which entrepreneurs who are more optimistic than the investor choose short-term debt, while entrepreneurs who share the investor’s belief choose long-term debt. In relation to our model, short-term debt may be interpreted as a bet in which the entrepreneur concedes cash flow rights in the low state in return for claims on the good state. The authors also provide empirical evidence suggesting that short-term debt is correlated with optimistic expectation errors of entrepreneurs. Sandroni and Squintani (2004) modify the Rothschild-Stiglitz insurance market model, to allow for consumers who are over-optimistic regarding their probability of an accident. They show that although in equilibrium these consumers are under-insured, compulsory insurance need not be Pareto-improving. Uthemann (2005) shows how to adapt a Hotelling-like model of competitive price discrimination due to Armstrong and Vickers (2001), in order to study competitive screening of consumers’ prior beliefs regarding the future value of a taste parameter.

Finally, this paper builds on our own previous work. Eliaz and Spiegel (2006) study optimal contract design with dynamically inconsistent agents. An agent type is his degree of naivete, modeled as his prior belief that his current preferences will change in the future. While the agent draws his belief from some distribution, the principal believes that the agent's preferences are sure to change. As in the present paper, the principal's objective in Eliaz and Spiegel (2006) is to screen agents according to their prior beliefs. However, time-inconsistency has important implications for the design of the optimal menu. In particular, sophisticated types who believe that their tastes will change with high probability are assigned a contract that serves as a perfect commitment device: it induces them to choose the action that maximizes their current utility.

Eliaz and Spiegel (2005) develop further the research agenda of mechanism design when agent types consist of their prior beliefs. The focus there is on bilateral speculation problems, where two agents hold different priors over an unverifiable state of nature, which affects the outcome of a game they are about to play. Eliaz and Spiegel (2005) define a notion of "constrained interim-efficient bets", characterize them and discuss their implementability in terms of the underlying game's payoff structure.

2 The model

A principal offers an agent the opportunity to choose an action from the set $[0, 1]$. The cost of providing an action a is $c(a)$, where $c(\cdot)$ is a continuous, non-decreasing function satisfying $c(0) = 0$. In order to have access to this set of actions, the agent must sign a contract with the principal one period beforehand. If the agent does not sign a contract with the principal, he chooses some outside option. We refer to the period in which a contract is signed as period 1, and to the period in which the action is chosen as period 2. A contract is a function $t : [0, 1] \rightarrow \mathbb{R}$ that specifies for every second-period action, a (possibly negative) transfer from the agent to the principal. The principal is perfectly able to monitor the agent's second-period action.

The agent has quasi-linear preferences over action-transfer pairs. We assume that his net utility in period 1 from the outside option is zero. However, his preferences over second period actions depend on the state of nature. There are two possible states: in state u the agent's preferences are represented by the continuous function $u : [0, 1] \rightarrow \mathbb{R}$, and in state v they are represented by the continuous function $v : [0, 1] \rightarrow \mathbb{R}$. We assume that there are always gains from trade ex-post - that is, $\max_a [u(a) - c(a)] \geq 0$ and $\max_a [v(a) - c(a)] \geq 0$.

The agent believes that state u occurs with probability θ . When $u(a) \geq v(a)$ for every a , it makes sense to refer to an agent with a higher θ as an agent with a higher degree of "optimism". Faced with a contract t , the agent's indirect utility from the contract is

$$\theta \max_{a \in [0,1]} [u(a) - t(a)] + (1 - \theta) \max_{a \in [0,1]} [v(a) - t(a)] \tag{1}$$

The principal believes that u occurs with probability p . The principal does not know the value of θ , and believes that it is distributed over $[0, 1]$ according to a continuous, strictly increasing *cdf* $F(\theta)$. Thus, the agent's "type" consists of his prior on u . The difference between the two parties' beliefs are purely due to differences in prior opinion: we assume that it is common knowledge that neither party is better informed about the state of nature.

The principal's objective is to maximize expected profits. By the revelation principle, a solution to his problem can be obtained via a direct revelation mechanism, in which agents are asked to report their type, and each reported type ϕ is assigned a contract $t_\phi : [0, 1] \rightarrow \mathbb{R}$. The optimal menu of contracts $\{t_\theta(a)\}_{\theta \in [0,1]}$ is given by the solution to the following maximization problem:

$$\max_{\{t_\theta(a)\}_{\theta \in [0,1]}} \int_0^1 \{p[t_\theta(a^u) - c(a^u)] + (1-p)[t_\theta(a^v) - c(a^v)]\} dF(\theta)$$

subject to the constraints,

$$\theta [u(a_\theta^u) - t_\theta(a_\theta^u)] + (1-\theta) [v(a_\theta^v) - t_\theta(a_\theta^v)] \geq 0 \quad (IR_\theta)$$

$$\theta [u(a_\theta^u) - t_\theta(a_\theta^u)] + (1-\theta) [v(a_\theta^v) - t_\theta(a_\theta^v)] \geq \theta [u(a_\phi^u) - t_\phi(a_\phi^u)] + (1-\theta) [v(a_\phi^v) - t_\phi(a_\phi^v)] \quad (IC_{\theta,\phi})$$

for all $\phi \in [0, 1]$, where

$$a_\theta^u \in \arg \max_{a \in [0,1]} [u(a) - t_\theta(a)] \quad (UR_\theta)$$

$$a_\theta^v \in \arg \max_{a \in [0,1]} [v(a) - t_\theta(a)] \quad (VR_\theta)$$

The first and second constraints are the standard individual rationality and incentive compatibility constraints. Condition IR_θ says that an agent of type θ is weakly better off with his assigned contract than with the default option. Condition $IC_{\theta,\phi}$ says that an agent of type θ cannot be better off by pretending to be of type ϕ and signing the contract assigned to that type.

The conditions UR_θ and VR_θ represent the fact that an agent's indirect utility from a contract is determined by the actions he expects to choose in the two states. If the realized state in period 2 is u (an event to which the agent assigns a probability of θ), then he will choose the optimal action for him according to the utility function u . This is represented by UR_θ . If, on the other hand, the state in period 2 is v (an event to which the agent assigns a probability of $1-\theta$), then he will choose the optimal action for him according to the utility function v . This is precisely the condition VR_θ .

It follows that any contract t can be identified with a pair of actions: a_θ^u and a_θ^v . The former action is consistent with u -maximization in the second period, while the latter action is consistent with v -maximization in the second period. Without loss of generality, we may assume that $t(a) = +\infty$ for every $a \notin \{a_\theta^v, a_\theta^u\}$.

The constraints IR_θ and $IC_{\theta,\phi}$ can be written more compactly by introducing the following notation. Let $D_\theta^u \equiv u(a_\theta^u) - t_\theta(a_\theta^u)$ and $D_\theta^v \equiv u(a_\theta^v) - t_\theta(a_\theta^v)$. Define $U(\phi, \theta)$ to be the expected payoff of a type θ agent who pretends to be of type ϕ - that is, $U(\phi, \theta) \equiv \theta D_\phi^u + (1 - \theta) D_\phi^v$. Then, IR_θ and $IC_{\theta,\phi}$ can be rewritten as $U(\theta, \theta) \geq 0$ and $U(\theta, \theta) \geq U(\phi, \theta)$ for all θ and ϕ . To further simplify the exposition, we shall sometimes use the following abbreviated notation. For $\omega = u, v$, let $t_\theta^\omega \equiv t_\theta(a_\theta^\omega)$, $c_\theta^\omega \equiv c(a_\theta^\omega)$, $a_{\max}^{\omega-c} \in \arg \max_a [\omega(a) - c(a)]$ and $\omega^* \equiv \omega(a_{\max}^{\omega-c})$.

3 A benchmark: common priors

Before we proceed to analyze the solution to the principal's problem in our model, it is instructive to consider a "benchmark" model, in which the principal and the agent agree on the probability of each state. There are two candidates for such a benchmark, depending on which side is believed to be better informed. Since we are interested in understanding how non-common priors affect the screening motives of a monopolist, we consider the benchmark to be a situation in which the principal believes that state u occurs with probability θ . As in our model, the value of θ is the agent's private information, and the principal believes that θ is distributed over $[0, 1]$ according to $F(\theta)$.

This benchmark essentially follows the framework used in the sequential-screening literature mentioned in the Introduction. This literature focuses on the case in which a higher type corresponds to a higher expected willingness to pay (for each action), or to a higher variance in the willingness to pay (for each action). In other words, the agents' types can be ordered according to first- or second-order stochastic dominance. The optimal menu in such an environment typically contains more than one contract. More importantly, the optimal menu is ex-post *inefficient* in the following sense: (i) there is a set of low types which are excluded from transacting with the monopolist, and (ii) all types (except for the highest one) choose sub-optimal actions.

In contrast, since our aim is to highlight the role of non-common priors in generating a non-degenerate optimal menu of contracts, we focus attention on a restricted domain of triples (u, v, c) , which turn out to imply that the optimal menu in the benchmark model consists of a single, ex-post efficient contract.

Definition 1 *We say that u and v satisfy the "crossing property" given a cost function c , if there exists a pair $(a_{\max}^{u-c}, a_{\max}^{v-c})$ that satisfy $u(a_{\max}^{u-c}) \geq u(a_{\max}^{v-c})$ and $v(a_{\max}^{v-c}) \geq v(a_{\max}^{u-c})$, with at least one strict inequality.*

The crossing property means that at the point in which gains from trade in a state reach a global maximum, the utility function of that state lies above the utility function of the other state. Note that for many natural cost functions, there is a large family of u and v functions that have the crossing property. For example, u and v may both be increasing

functions. Alternatively, u and v may be increasing and decreasing functions, respectively, provided that $v(0) = 0$.

Proposition 1 *Suppose that u and v satisfy the crossing property given c . Then, it is optimal for the principal to offer a single, ex-post efficient contract.*

Thus, under common priors, the principal cannot do better than to offer a contract that charges u^* if $a \in \arg \max(u - c)$ and v^* if $a \in \arg \max(v - c)$ (and an arbitrarily large amount for any other action). In particular, if $u^* = v^*$, the principal can simply “sell the project” to the agent in return for an up-front payment of u^* .

To conclude, the crossing property rules out discrimination and ex-post inefficiency in the common-prior benchmark. In contrast, in the next section we show that under non-common priors, the optimal menu typically discriminates among agent types. In environments which violate the crossing property, it would be difficult to disentangle the price discrimination that is due to speculation and that which arises from the considerations examined in the sequential-screening literature. For this reason, we shall henceforth assume that u and v satisfy the crossing property given c .

4 The optimal menu with non-common priors

4.1 Qualitative features of the menu

In this section we characterize the optimal menu of contracts, when the principal’s prior on u is p and the agent’s prior on u is θ , and the principal believes that θ is drawn from a distribution F on $[0, 1]$. We begin by illustrating the qualitative features of the optimal menu with a simple stylized example of a “backup agreement” between a supplier and a retailer.

Consider a retailer who buys from a supplier one unit of a good that is made up of two components, labeled U and V . The retailer has to determine the proportion of each component in the unit that he orders. Let $a \in [0, 1]$ denote the proportion of component U in the good. The retailer’s revenue as a function of a depends on the state of nature: in state u the revenue is given by $u(a) = a$, while in state v it is given by $v(a) = 1 - a$. The supplier has zero costs, and he believes that each state is equally likely. Therefore, he wishes to maximize the sum of transfers that he receives in the two states. The retailer, on the other hand, holds one of three possible prior beliefs: $\theta \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$.

Suppose the supplier knows that he is facing a retailer of type $\theta = \frac{1}{4}$. The supplier then wishes to maximize the sum of transfers, $t_{\frac{1}{4}}(a^u) + t_{\frac{1}{4}}(a^v)$, subject to two constraints: (i) given these transfers, the retailer would choose an action that maximizes his net payoff in each state, and (ii) the retailer weakly prefers to sign the contract. The first constraint represents

the UR and VR requirements, which imply:

$$-(a^u - a^v) \leq t_{\frac{1}{4}}(a^u) - t_{\frac{1}{4}}(a^v) \leq a^u - a^v$$

In order to relax this constraint as much as possible, the supplier should offer an ex-post efficient contract in which $a^u = 1$ and $a^v = 0$. The second constraint then implies:

$$\frac{1}{4}t_{\frac{1}{4}}(a^u) + \frac{3}{4}t_{\frac{1}{4}}(a^v) = 1$$

It follows that it is optimal for the supplier to offer the following two-part tariff:

$$t_{\frac{1}{4}}(a) = \frac{3}{4} + a$$

The interpretation is that the retailer buys one unit that contains only the V component at a price of $\frac{3}{4}$, and any amount of V can later be exchanged to U according to a one-to-one ratio.

Essentially the same argument holds when the supplier knows that he is facing a retailer of type $\theta = \frac{3}{4}$. In this case, it is optimal for the supplier to offer this retailer the contract:

$$t_{\frac{3}{4}}(a) = \frac{3}{4} + (1 - a)$$

The interpretation is that the retailer buys one unit that contains only the U component at a price of $\frac{3}{4}$, and any amount of U can later be exchanged to V according to a one-to-one ratio.

Finally, suppose the supplier knew he was facing a retailer who believes that each state is equally likely. By the $IR_{\frac{1}{2}}$ constraint, a retailer of this type is not willing to sacrifice a loss in one state in return for a smaller gain in the other state. Hence, the supplier cannot extract from this retailer more than a surplus of one in each state. This can be achieved with a contract that charges $t(a) = 1$ for all $a \in [0, 1]$.

Now assume the supplier cannot identify the retailer's type, and instead invites him to choose among the above three contracts: $\{t_{\frac{1}{4}}, t_{\frac{1}{2}}, t_{\frac{3}{4}}\}$. All contracts induce the retailer to choose $a = 1$ in state u and $a = 0$ in state v , hence all contracts are ex-post efficient. However, each retailer type weakly prefers his designated contract to the other contracts in the menu. For $\theta = \frac{1}{4}$, the expected payoff from both $t_{\frac{1}{4}}$ and $t_{\frac{1}{2}}$ is zero, while the expected payoff from $t_{\frac{3}{4}}$ is $-\frac{1}{2}$. Similarly, for $\theta = \frac{3}{4}$, the expected payoff from both $t_{\frac{3}{4}}$ and $t_{\frac{1}{2}}$ is zero, while the expected payoff from $t_{\frac{1}{4}}$ is $-\frac{1}{2}$. Finally, for $\theta = \frac{1}{2}$, the expected payoff from $t_{\frac{1}{2}}$ is zero, while the expected payoff from both $t_{\frac{1}{4}}$ and $t_{\frac{3}{4}}$ is $-\frac{1}{4}$. It follows that the above menu is optimal. The fact that all three types display the same behavior in period 2, while ending up making different payments, cannot be replicated by a standard model with a common prior.

This example illustrates some of the key features of the optimal menu summarized in the

Introduction. First, the optimal menu is the union of two separate menus: the optimal menu for types who assign a weakly higher prior on u , and the optimal menu for types who assign a weakly higher prior on v . The agent types belonging to one category exert no informational externality on the types belonging to the other category. Specifically, type $\frac{1}{4}$ strictly prefers his contract to the contract selected by type $\frac{3}{4}$, and vice versa.

The fact that each type is assigned the same contract he would be offered under perfect information is only due to the fact that in the example there is a single type $\theta > p$ and a single type $\theta < p$. As we shall see in the next subsection, when there are multiple agent types in each of these two categories, incentive constraints within each category may prevent the principal from attaining the perfect-information profit.

Second, the contracts signed by the types $\theta \neq p$ are essentially bets. Each of these types accepts a loss in the state which he deems less likely than the principal, in return for a gain in the other state. The payoffs that these types expect to obtain from their contracts are “speculative”, in the sense that they would not have signed these contracts had they shared the principal’s beliefs. Similarly, the principal would not have offered these contracts had he held the agents’ priors.

In contrast, the contract signed by the type $\theta = p$ involves no speculation: it guarantees a payoff of one, regardless of the agent’s prior. While the principal transacts with all agent types, he does not “bet” with the type who shares his prior belief. Note that the non-speculative contract in our example would be an optimal contract for *all* agent types in the benchmark model, in which it is commonly known that the agent’s type reflects the true probability that his second-period utility function is u .

Finally, the example demonstrates the following property of the optimal menu, which will turn out to be useful in characterizing the menu. An agent with a *higher* prior on u than the principal, who signs a speculative contract, strictly prefers a^u to a^v in state u , yet he is indifferent between a^u and a^v in state v . Similarly, an agent with a *lower* prior on u than the principal, who signs a speculative contract, strictly prefers a^v to a^u in state v , yet he is indifferent between a^u and a^v in state u . Put differently, the VR constraint is binding for speculative agents who are more confident than the principal about state u , while the UR constraint is binding for speculative agents who are more confident than the principal about state v .

4.2 Characterization

Our goal in this subsection is to provide a complete characterization of the optimal menu of contracts. In the process of reaching this goal, we shall uncover some interesting properties of this menu.

In the example of the previous subsection, the agent types who held different beliefs than the principal were offered “speculative contracts”, which they would never accept had they shared the principal’s prior. Formally, we define such contracts as follows.

Definition 2 *A contract t is speculative if*

$$p[u(a^u) - t(a^u)] + (1 - p)[v(a^v) - t(a^v)] < 0 \quad (2)$$

The optimal menu in our example included both speculative and non-speculative contracts. The non-speculative contract was the optimal contract to offer to an agent who shares the principal's beliefs. Our first result establishes that this property also extends to the case with a continuum of agent types.

Lemma 1 *If the optimal menu includes a non-speculative contract, then that contract must be the first-best contract for an agent of type $\theta = p$.*

Lemma 1 implies that w.l.o.g. the optimal menu includes at most one non-speculative contract. This contract is ex-post efficient and extracts from the agent u^* in state u and v^* in state v . Hence, an agent who signs this contract expects to obtain his reservation utility regardless of his prior.

The intuition for this result is as follows. The profit that the principal expects to obtain from any contract t is $p[t(a^u) - c(a^u)] + (1 - p)[t(a^v) - c(a^v)]$, where $a^u \in \arg \max_a [u(a) - t(a)]$ and a^v is similarly defined. By definition, if t is non-speculative, then this expected profit is at most $p[u(a^u) - c(a^u)] + (1 - p)[v(a^v) - c(a^v)]$. This expression is maximized at $a^u \in \arg \max(u - c)$ and $a^v \in \arg \max(v - c)$. Therefore, the following contract t^* is an optimal non-speculative contract: $t^*(a) = u(a)$ if $a \in \arg \max(u - c)$, $t^*(a) = v(a)$ if $a \in \arg \max(v - c)$ and $t^*(a) = \infty$ for every other a .

By the crossing property, t^* satisfies the UR and VR constraints of all agent types. If the principal's menu included non-speculative contracts that generate a lower expected profit, the principal could omit them from the menu, without causing agent types to switch from speculative contracts to t^* . The reason is that t^* yields an expected payoff of zero to all types, hence the constraint that forbids an agent type to prefer t^* to his designated contract is indistinguishable from his IR constraint.

The optimal menu in our example contained both speculative and non-speculative contracts. A natural question that arises is whether this is a general feature of an optimal menu. By IR_p , type $\theta = p$ cannot be assigned a speculative contract. The question is, does an optimal menu necessarily include a speculative contract?

Lemma 2 *If $0 < p < 1$, the optimal menu contains at least one speculative contract. If $p = 0$, the optimal menu includes at least one speculative contract if and only if $u(a) > v(a)$ for some a . If $p = 1$, the optimal menu includes at least one speculative contract if and only if $u(a) < v(a)$ for some a .*

To understand the intuition for this result, consider how a speculative contract exploits the disagreement between the principal and the agent. For the sake of the argument, let $p = 0$. If $u(a) > v(a)$ for some a , then the principal can offer a contract that charges $u(a) - \varepsilon$ for choosing $a \in \arg \max[u(a) - v(a)]$, but charges $v^* + \max[u(a) - v(a)] - \varepsilon$ for choosing any action that maximizes $v(a) - c(a)$ (and an infinite amount for any other action). Any agent who accepts this contract gains a surplus of ε in state u , but loses a surplus of $\max[u(a) - v(a)] - \varepsilon$ in state v . High- θ types weigh the gain more than the loss, and hence prefer this contract to a non-speculative contract that leaves them with zero expected payoff. For the principal, who believes that state v is certain to occur, this contract generates more profit than the highest total surplus in state v . Hence, he prefers this contract to the non-speculative contract. If, however, $v(a) \geq u(a)$ for all a , the principal cannot entice optimistic types by offering them an “imaginary win” in state u in return for a “real fine” in state v .

If the optimal menu contains both speculative and non-speculative contracts, which agent types are assigned the former and which are assigned the latter? Our next result provides a simple answer.

Lemma 3 *There exists a pair of types, $\underline{\theta} \in [0, p)$ and $\bar{\theta} \in (p, 1]$ such that: (i) t_θ is non-speculative for every $\theta \in (\underline{\theta}, \bar{\theta})$, and if $\underline{\theta} = 0$ then t_0 is also non-speculative, and similarly if $\bar{\theta} = 1$ then t_1 is non-speculative; (ii) t_θ is speculative for every $\theta < \underline{\theta}$; (iii) t_θ is speculative for every $\theta > \bar{\theta}$.*

This result stems from ordinary single-crossing arguments. The contract offered to $\theta = p$ must be non-speculative, by IR_p . Suppose that the optimal menu assigns a speculative contract to some type $\theta > p$, but assigns a non-speculative contract to a type $\phi > \theta$. By Definition 2,

$$\begin{aligned} (p)D_\phi^u + (1-p)D_\phi^v &\geq 0 \\ (p)D_\theta^u + (1-p)D_\theta^v &< 0 \end{aligned}$$

Because $\theta > p$, these inequalities imply that $IC_{\phi, \theta}$ is violated. A similar argument shows that if type $\theta' < p$ is assigned a speculative contract, then every type $\phi' < \theta'$ must also be assigned such a speculative contract.

Lemma 3 implies that the “threshold” types, $\bar{\theta}$ and $\underline{\theta}$, should be indifferent between a speculative contract and a non-speculative contract. Since, by Lemma 1, a non-speculative contract gives a zero expected payoff to all types, the speculative contract offered to $\bar{\theta}$ and $\underline{\theta}$ should also give these types an expected payoff of zero. This implies the following result.

Corollary 1 *$IR_{\underline{\theta}}$ and $IR_{\bar{\theta}}$ must be binding.*

Recall that in the example of the previous subsection, the principal chose to speculate with types whose priors lie on both sides of his own prior. In some situations, however, the principal chooses to speculate only with agents whose priors lie on one side of p .

Proposition 2 (“No speculation with pessimists”) *If $u(a) \geq v(a)$ for all a , then $\underline{\theta} = 0$. Similarly, if $v(a) \geq u(a)$ for all a , then $\bar{\theta} = 1$.*

The meaning of this result is that if payoffs in one state are always (weakly) higher than in another state - such that θ may be viewed as the agent’s “degree of optimism” - then the principal chooses not to speculate with agent types who are more pessimistic than he is. For example, suppose the principal and the agent are a manager and a sales agent, such that u and v represent the agent’s profits from generating a sales volume of a . Note that if profits are positively correlated with ability, then in a situation where u is above v , θ may be interpreted as the agent’s degree of “overconfidence”. In this case, Proposition 2 implies that the manager would speculate only with overconfident agents.

Another distinctive feature of the example of the previous subsection is that the speculative contract designed for type $\theta > p$ is designed independently of the contract designed for $\theta < p$. Our next result establishes that this is a general property of the optimal menu.

Lemma 4 *The optimal menu is the union of two sets of menus: (i) the optimal menu for the distribution F conditional on the restricted support $[p, 1]$, and (ii) the optimal menu for the distribution F conditional on the restricted support $[0, p]$.*

The essence of Lemma 4 is that the speculative contracts offered to types $\theta < p$ do not exert an informational externality on types $\theta > p$ who speculate with the principal, and vice versa. Hence, the only incentive constraints that the principal needs to worry about are those that prevent agents whose prior is on one side of p to exaggerate the proximity of their belief to the principal’s. The proof of this result is based on the definition of speculative contracts and a single-crossing argument.

Lemma 4 simplifies the derivation of the optimal menu, in that it breaks it down into three separate problems: (i) solving for the optimal menu for $\theta \geq \bar{\theta}$, (ii) solving for the optimal menu for $\theta \leq \underline{\theta}$, and (iii) solving for the non-speculative contract for $\underline{\theta} \leq \theta \leq \bar{\theta}$. By Lemma 1, the non-speculative contract is immediate. The question is, how does one design the optimal menus of speculative contracts? In standard models of price discrimination, a menu consists of *pairs* of numbers (e.g., quantity and price). In contrast, a menu in our model consists (w.l.o.g.) of *two pairs* of numbers: (a_θ^u, t_θ^u) and (a_θ^v, t_θ^v) . Therefore, in order to be able to apply standard tools of optimal price discrimination, we need to further simplify our problem. The first step in this simplification is to show that w.l.o.g. we may fix a_θ^v for $\theta \geq p$ and a_θ^u for $\theta \leq p$.

Lemma 5 *Without loss of generality, we may restrict attention to an optimal menu in which $a_\theta^v \in \arg \max_a [v(a) - c(a)]$ for $\theta \geq p$, and $a_\theta^u \in \arg \max_a [u(a) - c(a)]$ for $\theta \leq p$.*

Lemma 5 implies that the agent chooses the ex-post efficient action in the state he deems less likely than the principal. The intuition is that even if ex-post inefficiency may be required for the sake of increasing speculative gains, the principal prefers this departure to occur in the state that he deems less likely.

The expected payoff of an agent of type $\theta > \bar{\theta}$, who signs a speculative contract, may be written as $\theta(D_\theta^u - D_\theta^v) + D_\theta^v$. Thus, it is as if the agent gives the principal an up-front payment of $-D_\theta^v$, and if state u is realized, the principal pays the agent the amount $D_\theta^u - D_\theta^v$. We may therefore interpret the difference $D_\theta^u - D_\theta^v$ as the “speculative gain” of type θ from the contract that he accepts. Denote $q(\theta) \equiv D_\theta^u - D_\theta^v$. By Definition 2, the speculative contract signed by θ gives a strictly negative surplus to an agent with a prior of p . Thus, in order for type $\theta > p$ to accept it, it must be the case that $q(\theta) > 0$. Applying the same argument to types $\theta \leq \underline{\theta}$, we obtain the following result.³

Lemma 6 *$q(\theta) > 0$ for all $\theta \geq \bar{\theta}$ and $q(\theta) < 0$ for all $\theta \leq \underline{\theta}$.*

A standard technique in the mechanism design literature is to transform the incentive-compatibility constraints into an integral representation of $U(\theta, \theta)$ (the expected utility of type θ from truthfully reporting his type). Lemma 6 is instrumental in adapting this technique to our framework. Rewrite $U(\theta, \theta)$ as follows:

$$U(\theta, \theta) = \begin{cases} \theta q(\theta) + D_\theta^v & \text{if } \theta \geq p \\ (1 - \theta)[-q(\theta)] + D_\theta^u & \text{if } \theta < p \end{cases} \quad (3)$$

By Lemma 6, we may apply standard arguments to obtain the following result.

Lemma 7 *Assume that the optimal menu assigns a speculative contract to type θ . If this contract satisfies $IC_{\theta, \phi}$ for all θ and ϕ , then*

$$U(\theta, \theta) = \begin{cases} \int_{\bar{\theta}}^{\theta} q(x) dx & \text{if } \theta \geq \bar{\theta} \\ \int_{1-\underline{\theta}}^{1-\theta} [-q(x)] dx & \text{if } \theta \leq \underline{\theta} \end{cases} \quad (4)$$

³For a formal proof of this Lemma, see Observation 1, which is found in the proof of Lemma 3 in the Appendix.

Recall that in our example of the previous sub-section, type $\theta > p$ was indifferent between a_θ^v and a_θ^u in state v , while type $\theta < p$ was indifferent between a_θ^v and a_θ^u in state u . That is, agent types who signed speculative contracts were indifferent between the two actions in the state they deemed less likely than the principal. Our next result establishes that this is a general property of the optimal menu.

Lemma 8 *The VR_θ constraint is binding for all types $\theta \geq \bar{\theta}$, while the UR_θ constraint is binding for all types $\theta \leq \underline{\theta}$.*

Lemmas 5 and 8 reduce the number of decision variables per type to only *two*: (a_θ^v, t_θ^v) for $\theta \leq \underline{\theta}$ and (a_θ^u, t_θ^u) for $\theta \geq \bar{\theta}$. Thus, we are now ready to apply standard mechanism-design tools to derive the optimal speculative menu. Let $z = 1 - x$ and $G(z) = 1 - F(1 - x)$. Define

$$\psi_F(x) = x - \frac{1 - F(x)}{f(x)} \quad (5)$$

and define $\psi_G(z)$ accordingly.

Proposition 3 *If F satisfies the monotone hazard rate condition, then a speculative contract intended for a type $\theta > p$ induces the actions,*

$$\begin{aligned} a_\theta^v &\in \arg \max_a [v(a) - c(a)] \\ a_\theta^u &\in \arg \max_a [(\psi_F(\theta) - p) \cdot (u(a) - v(a)) + p(u(a) - c(a))] \end{aligned} \quad (6)$$

If G satisfies the monotone hazard rate condition,⁴ then a speculative contract intended for a type $\theta < p$ induces the actions:

$$\begin{aligned} a_\theta^u &\in \arg \max_a [u(a) - c(a)] \\ a_\theta^v &\in \arg \max_a [(\psi_G(1 - \theta) - (1 - p)) \cdot (v(a) - u(a)) + (1 - p) \cdot (v(a) - c(a))] \end{aligned} \quad (7)$$

The speculative contracts in the optimal menu may induce an *inefficient* action in the state which the agent deems more likely than the principal. To interpret this action, consider the case of $p = 0$. The principal would like the agent to choose the efficient action in state v (the state he believes with certainty), but he could exploit his disagreement with the agent to charge more than v^* for this action. Each contract in the principal's menu has three components: an action a^u , an up-front payment t_θ , and a changing fee of τ_θ . An agent who

⁴For instance, when F is the uniform distribution, both F and G satisfy the monotone hazard rate property.

signs a contract $(a^u, t_\theta, \tau_\theta)$ pays t_θ and commits to choosing the action a^u (which is the same for all speculative contracts because $p = 0$). However, if in period 2 the agent switches to $\arg \max_a (v(a) - c(a))$, he has to pay τ_θ . The principal believes that the agent will pay the fee for sure, while the agent believes that this will happen with probability θ . The principal's objective is to have $t_\theta + \tau_\theta > v^*$, and as high as possible. To meet this objective, the principal should choose a^u such that the agent will agree to pay as much as possible in state v for switching from a^u to $\arg \max_a (v(a) - c(a))$. Hence, $a^u \in \arg \max_a [u(a) - v(a)]$.

This intuition extends to the case of $p \in (0, 1)$. This case is a little more complicated. First, we need to distinguish between optimistic and pessimistic types. The former pay the "changing fee" in state v , while the latter pay this fee in state u . Second, since the principal assigns positive probability to both states, the action a^u for the optimistic types maximizes a weighted average of $u - v$ and $u - c$. Similarly, a^v for the pessimistic types maximizes a weighted average of $v - u$ and $v - c$.

Notice that if the principal observed the agent's type, then in order to derive a_θ^u in the first-best contract for a type $\theta > p$, we would simply need to replace $\psi_F(\theta)$ with θ in (6). Likewise, we would replace $\psi_G(1 - \theta)$ with $1 - \theta$ in (7) when deriving a_θ^v in the first-best contract for a type $\theta < p$.

To complete our characterization of the optimal menu, it remains to derive the cutoffs, $\bar{\theta}$ and $\underline{\theta}$. Recall that by Proposition 2, if $u(a) \geq v(a)$ (respectively, $v(a) \geq u(a)$) for all a , then the principal weakly prefers not to speculate with any of the types below p (above p), in which case $\underline{\theta} = 0$ ($\bar{\theta} = 1$). The following proposition characterizes the cutoffs for the case in which $u(a) > v(a)$ for some a and $v(a) > u(a)$ for some a .

Proposition 4 (i) *Suppose that F satisfies the monotone hazard rate condition, and $u(a) > v(a)$ at some a . Then, $\bar{\theta}$ is the unique solution of $\psi_F(\theta) = p$.*

(ii) *Suppose that G satisfies the monotone hazard rate condition, and $v(a) > u(a)$ at some a . Then, $\underline{\theta}$ is the unique solution of $\psi_G(1 - \theta) = 1 - p$.*

Our next collection of results addresses the question of ex-post efficiency of the optimal menu. We say that the optimal menu is ex-post efficient if $a_\theta^u \in \arg \max(u - c)$ and $a_\theta^v \in \arg \max(v - c)$ for every type θ .

Proposition 5 *If the optimal menu is ex-post efficient, then there exists such a menu with at most one speculative contract for types higher than p , and at most one such contract for types lower than p .*

If the optimal menu is ex-post efficient, then w.l.o.g. all types can be made to choose precisely the same action in each state ($a^u \in \arg \max(u - c)$ and $a^v \in \arg \max(v - c)$). But

since each optimistic type is indifferent between a^u and a^v in state v , while each pessimistic type is indifferent between these two actions in state u (recall Lemma 8), the absolute value of the “speculative gain” $|q(\theta)|$ is constant for *all* types. Hence, the principal can only discriminate between types $\theta > p$ and types $\theta < p$, but he cannot discriminate among the types within each of these categories.

The following result provides a sufficient condition for ex-post efficiency.

Proposition 6 *If $\arg \max(u - c) = \arg \max(u - v)$ and $\arg \max(v - c) = \arg \max(v - u)$, then the optimal menu is ex-post efficient .*

The intuition for this result is simple. Consider the speculative contracts that the principal designs for types $\theta > p$. By Lemma 5, a_θ^v is ex-post efficient. In our discussion below Proposition 3, we remarked that when designing a_θ^u , the principal attempts to strike a compromise between maximizing $u - c$ and maximizing $u - v$. When the same action maximizes both $u - c$ and $u - v$, the principal can set a^u to be ex-post efficient without having to worry about this trade-off. Note that the sufficient condition identified in the Proposition is satisfied in our example of the previous sub-section.

5 Examples

In this section we illustrate some features of the optimal menu using a collection of simple applications of our framework. Given the simplicity of our model, the examples are highly stylized and are not meant to serve as descriptive models of the concrete economic environments referred to. However, we believe they illuminate contractual arrangements that we observe in reality. Throughout this subsection, we assume $F(\theta) = \theta$ and $p = \frac{1}{2}$.

5.1 Ownership vs. fixed wages

A manager contemplates hiring a sales agents. Assume that the volume of sales can take any value in the interval $[0, 1]$. The manager incurs no costs except for his payments to the agent (i.e., $c(a) = 0$ for all a). The cost that the agent incurs in generating a sales volume of a depends on the state of nature. In state u , it is not costly to generate higher sales, whereas in state v , the cost of generating sales exactly offsets the revenue. Thus, $u(a) = a$ and $v(a) = 0$ for every $a \in [0, 1]$. The agent’s reservation value, normalized to zero, is interpreted as the highest wage he can earn if he declines all of the manager’s offers.

Our results in the previous section greatly simplify the task of deriving the optimal menu. By Proposition 2, the manager will speculate only with “optimistic” sales agents. By Proposition 6, the optimal contract is ex-post efficient. Therefore, by Proposition 5, the menu will consist of one non-speculative contract (denoted t_{NS}) and one speculative contract (denoted t_S). All types lower than $\bar{\theta}$ sign t_{NS} , while all higher types sign t_S .

By Lemma 1, the non-speculative contract extracts the entire surplus of the agent in each of the states. One such contract is $t_{NS}(a) = a$. This contract means that the agent receives no share in the sales that he generates, and he is compensated by a fixed wage which is equal to his reservation value.

To construct t_S , we first need to find the threshold type $\bar{\theta}$. This type satisfies that the marginal gain in expected profit from speculating with him is equal to the marginal loss from not speculating with him, i.e.,

$$p \max_a u(a) = \psi_F(\bar{\theta})[u(a^u) - v(a^u)] + pv(a^u) \quad (8)$$

By Proposition 3, $a^u = 1$ and $a^v = 0$. This means that (8) may be written as

$$\frac{1}{2} \cdot 1 = (2\bar{\theta} - 1) \cdot 1 + \frac{1}{2} \cdot 0$$

Therefore, $\bar{\theta} = \frac{3}{4}$.

The pair of transfers, $t_S(0)$ and $t_S(1)$, can now be derived using the following observations. By Lemma 8, the VR constraint is binding for all agents who sign t_S :

$$-t_S(0) = -t_S(1) \quad (9)$$

By Corollary 1, type $\frac{3}{4}$ is indifferent between signing a speculative contract and signing a non-speculative contract:

$$\frac{3}{4}[1 - t_S(1)] + \frac{1}{4}[-t_S(0)] = 0 \quad (10)$$

Taken together, (9) and (10) imply that $t_S(0) = t_S(1) = \frac{3}{4}$.

These transfers may be implemented by a contract, $t_S(a) = \frac{3}{4}$ for every a . This contract essentially “sells the project” to the sales agent. Thus, the optimal contract consists of two extreme risk-sharing schemes. In one scheme, the manager bears all the risk and the sales person receives a fixed wage. In another scheme, the sales agent obtains full ownership of the project and bears all the risk.

Both contracts provide precisely the same incentives in the second period. Under both contracts, the agent will generate a sales volume $a^u = 1$ in state u and $a^v = 0$ in state v . Thus, the multiplicity of contracts in the menu has nothing to do with second-period incentives, or with the agent’s risk attitudes. Instead, the agent’s selection of a payment scheme reveals his degree of optimism.

This example is special, in the sense that t_S leaves the manager totally insured. In this respect, the term “speculative contract” is perhaps a misnomer, because the manager does not bear any risk under t_S . This feature is due to the assumption that v is flat (rather than, say, a strictly decreasing function). In the examples that follow, the optimal menu will contain speculative contracts that impose risk on both parties.

5.2 Unlimited consumption vs. variable rates

A commonly observed menu of pricing schemes offers unlimited consumption at a fixed fee, side by side with a variable-rate scheme that charges according to consumption. Menus of this type are offered by telecommunication companies, where firms often offer a choice between “unlimited calling plans” and plans that condition per-minute rate to the amount of minutes used. Similarly, DVD rental stores offer “unlimited plans” as well as “limited plans”. Finally, as in the example presented in the Introduction, restaurants sometimes offer diners a choice between an “all you can eat” buffet and a selection of dishes “a la carte”.

We propose to interpret such a menu as a tool for screening consumers according to their prior beliefs regarding their future tastes. This interpretation fits situations in which the agent is unsure of his future satiation level: would he desire a large amount of airtime to communicate with his partners on the mobile phone, or will short conversations suffice? how much idle time will he have for watching rental movies? how much would he need to eat in order to satisfy his appetite?

In these situations, a variable-rate contract may be viewed as a bet between the principal and the agent, where the agent “wins” if he manages to consume only a small amount and the principal “wins” if the agent ends up consuming a large amount. In contrast, an unlimited consumption contract has no speculative component, because the payment the agent makes is independent of his level of consumption. Consumers who believe that their satiation level is likely to be low would prefer a speculative, variable-rate contract. On the other hand, consumers who believe that their satiation level is likely to be high would opt for the non-speculative, unlimited-consumption contract.

We illustrate this idea with a simple example. A consumer is about to purchase a calling plan from a monopolistic mobile phone company. Being a newcomer to the mobile phone market, the consumer does not know if his required amount of airtime will be high or low. Formally, let a denote the proportion of monthly minutes that the consumer uses (where $a = 1$ means that a consumer uses every available minute). The costs to the phone company are given by $c(a) = \frac{1}{4}a$. Let $u(a)$ and $v(a)$ represent the consumer’s willingness to pay for a minutes of airtime in the low and high state respectively. Specifically, let $v(a) = a$ and

$$u(a) = \begin{cases} 2a & \text{for } 0 \leq a < \frac{1}{2} \\ 1 & \text{for } \frac{1}{2} \leq a < 1 \end{cases}$$

Note that from the mobile phone company’s point of view, a consumer of type $\theta > \frac{1}{2}$ ($\theta < \frac{1}{2}$) underestimates (overestimates) his required amount of airtime. Because $u(a) \geq v(a)$ for every a , Proposition 2 implies that the phone company will offer a non-speculative contract to types below $\bar{\theta} > \frac{1}{2}$. Such a contract can be implemented by a calling plan that offers unlimited calls for a monthly fee of 1.

To derive the set of speculative contracts, we first note that by Proposition 3, for all

$\theta \geq \bar{\theta}$, $a_\theta^v = 1$ and

$$a_\theta^u \in \arg \max_{a \in [0,1]} [(2\theta - 1)u(a) + (\frac{5}{4} - 2\theta)a]$$

Solving this optimization problem yields $a_\theta^u = \frac{1}{2}$. Since $a_\theta^u \in \arg \max_a [u(a) - c(a)]$ and $a_\theta^v \in \arg \max_a [v(a) - c(a)]$, it follows that the optimal menu is ex-post efficient. By Proposition 5, the optimal menu contains a single speculative contract t^* .

Recall that VR_θ is binding for every $\theta \geq \bar{\theta}$ and that IR_θ is binding for $\bar{\theta}$. By Proposition 4, $\bar{\theta} = \frac{3}{4}$. Therefore:

$$\begin{aligned} \frac{3}{4}[u(\frac{1}{2}) - v(1) + t^*(1) - t^*(\frac{1}{2})] + v(1) - t^*(1) &= 0 \\ t^*(1) - t^*(\frac{1}{2}) &= v(1) - v(\frac{1}{2}) \end{aligned}$$

These equations yield $t^*(1) = \frac{11}{8}$ and $t^*(\frac{1}{2}) = \frac{7}{8}$. The speculative contract t^* can thus be implemented by the following three-part tariff. For a flat fee of $\frac{7}{8}$, the consumer receives $a = \frac{1}{2}$ at a marginal price of zero. For any additional units, the consumer has to pay a marginal price of 1 (i.e., $t(a) = \frac{3}{8} + a$ for $a > \frac{1}{2}$).

Thus, we have interpreted the coexistence of variable-rate and unlimited consumption contracts as a consequence of the principal's attempt to screen the agent's probabilistic assessment of his satiation point. The unlimited consumption contract is a non-speculative contract, whereas the variable-rate contract is a speculative contract directed at consumers who significantly underestimate (in a probabilistic sense) their satiation level of consumption.

Applied to the DVD rental example, our characterization of the optimal menu implies that unlimited rental plans are essentially non-speculative contracts, while limited plans are speculative contracts aimed at consumers who underestimate their future amount of idle time.

Our interpretation is perhaps surprising in the context of the restaurant example. A priori, one might think that the "all you can eat" deal should be viewed as a speculative contract, targeted at diners who overestimate their satiation level. However, according to our interpretation, this deal is a non-speculative contract, while the "a la carte" deal is a speculative contract, targeted at customers who enter the restaurant believing they will have only an entree and a main course, but find themselves browsing the dessert menu at the end of the meal .

5.3 Speculative contracts and inefficiency

In many situations a buyer contracts with a supplier prior to knowing the exact specifications of the product or service he requires, as these depend on a future state of nature. Bilateral contracts of this nature have been widely studied in the literature. Most often, the focus has been on the hold-up problem that may arise and the various ways in which it could be remedied (see Tirole (1999) and the references therein). An important aspect of these contingent contracts that has been largely overlooked, is their speculative role in allowing

parties with heterogeneous beliefs to bet on the future (see Bazerman and Gillespie (1999)). In this sub-section we analyze a bilateral contracting example of this type. We also employ this example to demonstrate the subtle effect of speculative contracts on ex-post efficiency.

A seller can provide (at zero cost) a product at any variety $a \in [0, 1]$ demanded by a buyer. The ideal variety for the buyer depends on the state of nature. Assume that u and v take the following functional forms:

$$\begin{aligned} u(a) &= 2 - \left(a - \frac{1}{2}\right)^2 \\ v(a) &= 2 - \left(a - \frac{3}{4}\right)^2 \end{aligned}$$

That is, the buyer's ideal variety is $a = \frac{1}{2}$ in state u and $a = \frac{3}{4}$ in state v . For expositional simplicity, we shall focus on the contracts that the seller designs for types $\theta \in [\frac{1}{2}, 1]$.

The non-speculative contract for types $\frac{1}{2} \leq \theta < \bar{\theta}$ may take the simple form of a flat payment of 2, which is independent of a . Let us turn to the speculative contracts for types $\theta \geq \bar{\theta}$. We first note that for each of these types, a_θ^v can be set to $\frac{3}{4}$. By Proposition 3, to compute a_θ^u for $\theta \geq \bar{\theta}$ we need to solve

$$\max_{a \in [0, 1]} \left\{ \left[(2\theta - 1) - \frac{1}{2} \right] \left[\left(a - \frac{3}{4}\right)^2 - \left(a - \frac{1}{2}\right)^2 \right] + \frac{1}{2} \left[2 - \left(a - \frac{1}{2}\right)^2 \right] \right\}$$

yielding $a_\theta^u = \frac{5}{4} - \theta$. Finally, by Proposition 4, $\bar{\theta} = \frac{3}{4}$.

This result has two noteworthy features. First, the non-speculative contract specifies a flat payment, whereas the speculative contract specifies a contingent pricing schedule. Buyers who roughly share the seller's belief choose a contract that gives them the freedom to choose the exact variety of the product only after they learn their ideal point. In contrast, buyers who sharply disagree with the seller are willing to commit to the variety they will choose and pay a fine in case they change their mind in period 2.

Second, note that $a_\theta^u \leq \arg \max u$ for all $\theta \geq \bar{\theta}$. Since a_θ^u decreases with θ , the distance between a_θ^u and $\arg \max u$ increases with θ (in the range $\theta > \frac{3}{4}$). Thus, as the parties' beliefs become more polarized, the contract they sign becomes more inefficient ex-post (in state u). In other words, the more speculative the contract, the more inefficient the action that it induces in state u .

The latter aspect of our example has a further implication. If the seller could observe the buyer's type θ , he would assign to any buyer type $\theta > \frac{1}{2}$ a speculative contract that induces $a_\theta^v = \frac{3}{4}$ and $a_\theta^u = \arg \max [(\theta - p) (u - v) + p u] = \frac{3}{4} - \frac{\theta}{2}$. Compare this with our result that when the seller does not observe θ , $a_\theta^u = \frac{1}{2}$ for $\theta \in [\frac{1}{2}, \frac{3}{4})$ and $a_\theta^u = \frac{5}{4} - \theta$ for $\theta > \frac{3}{4}$. It is easy to see that the outcome is "less inefficient" ex-post when the seller does not observe the buyer's type. Thus, if a social planner, who wishes to maximize social surplus (according to his own prior beliefs), had to choose between an environment in which the seller observes the

buyer's prior and an environment in which the buyer's prior is his private information, he would prefer the latter environment.

These welfare implications are not general, but a consequence of certain features of the payoff structure: (i) $v(a) \equiv u(a - d)$, where d is the distance between the ideal points in the two states; (ii) u is concave - i.e., as the distance from the ideal point becomes larger, the marginal disutility from steering away from it increases; (iii) the Arrow-Pratt coefficient $-u''/u'$ increases with a (in the relevant range, in which a falls below the ideal point). The proof is elementary, involving first- and second-order conditions of the objective function $(\theta - p)(u - v) + pu$ when the agent's prior is known, and $(\psi_F(\theta) - p)(u - v) + pu$ when the prior is private information.

Note that in contrast to previous examples in the paper, the optimal menu in this subsection displays fine discrimination among optimistic types. Specifically, there is a continuum of speculative contracts. To characterize these contracts for $\theta > \frac{3}{4}$, recall that VR_θ is binding for these types. Therefore,

$$t(a_\theta^v) - t(a_\theta^u) = \left(\frac{1}{2} - \theta\right)^2 \quad (11)$$

This implies that for all $\theta \geq \frac{3}{4}$,

$$q(\theta) = \frac{\theta}{2} - \frac{5}{16}$$

Hence, by (18),

$$t(a_\theta^v) = \theta\left(\frac{\theta}{2} - \frac{5}{16}\right) + 2 - \int_{\frac{3}{4}}^{\theta} \left(\frac{x}{2} - \frac{5}{16}\right) dx = \frac{\theta^2}{4} + \frac{61}{32}$$

Substituting this into (11) we obtain $t(a_\theta^u) = t(a_\theta^v) - \left(\frac{1}{2} - \theta\right)^2$. This means that the higher the buyer's prior on u , the lower the payment he makes in this state, and the higher the payment he makes in v .

6 Conclusion

We have argued that menus of contingent contracts in principal-agent environments can be explained as a consequence of the principal's attempt to screen the agent's prior belief, when there are non-common priors. Whereas standard accounts highlight the role of contingent contracts in providing incentives for the agent, we interpret them as bets. Indeed, in many of our examples the agent's actions are independent of the contract he selects from the menu, and therefore incentive provision *cannot* be the explanation for price discrimination.

We do not take a stand as to which of the two parties, if any, holds the correct prior. Their differences in beliefs may or may not be a result of inherent biases such as over-optimism. For example, in the manager-worker example of Subsection 5.1, each party may be over-optimistic or over-pessimistic regarding the agent's future cost of effort, and our analysis does not require us to make any judgment in this regard. Even if it is natural to assume that the agent will tend to be over-optimistic relative to the principal, it does not follow that we

have to assume that the principal is right.

In some applications, however - e.g., when the principal is a firm and the agent is a consumer - one might want to interpret the principal's prior p as correct, and the agent's prior $\theta \neq p$ as resulting from a psychological bias such as over-optimism. In this case, in order to be consistent with the model, it must be assumed that the consumer does not believe that his belief is biased. Otherwise, he would regard the menu as a signal of the principal's prior, and he would update his beliefs. Also, such an interpretation would have to be confronted with the empirical question of whether consumers make systematic errors in anticipating their future tastes (see Miravete (2003), for example).

References

- [1] Armstrong, M. (1996), "Nonlinear Pricing with Imperfectly Informed Consumers," mimeo, Nuffield College.
- [2] Armstrong, M. and J. Vickers (2001), "Competitive Price Discrimination," *Rand Journal of Economics*, 32, 579-605.
- [3] Baron, D.P. and D. Besanko (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1(3), 267-302.
- [4] Bazerman, M. H., and J. J. Gillespie (1999), "Betting on the Future: The Virtues of Contingent Contracts," *Harvard Business Review*, 155-160.
- [5] Courty, P. and H. Li (2000), "Sequential Screening," *Review of Economic Studies*, 67(4), 697-718.
- [6] Eliaz, K. and R. Spiegler (2005), "A Mechanism Design Approach to Speculative Trade," mimeo, New York University and Tel-Aviv University.
- [7] Eliaz, K. and R. Spiegler (2006), "Contracting with Diversely Naive Agents," *Review of Economic Studies*, 73(3), 689-714.
- [8] Fang, H. and G. Moscarini (2005), "Morale Hazard," *Journal of Monetary Economics*, 52, 749-777.
- [9] Krishna, V. (2002), *Auction Theory*, Academic Press, San Diego.
- [10] Landier, A. and D. Thesmar (2004), "Financial Contracting with Optimistic Entrepreneurs," mimeo, University of Chicago GSB.
- [11] Miravete, E.J. (2003), "Choosing the Wrong Calling Plan: Ignorance and Learning," *American Economic Review*, 93(1), 297-310.

- [12] Sandroni A. and F. Squintani (2004), “The Over-Confidence Problem in Insurance Markets,” mimeo, UCL.
- [13] Tirole, J. (1999), “Incomplete Contracts: Where Do We Stand?,” *Econometrica*, 67(4), 741-781.
- [14] Uthemann, A. (2005), “Competitive Screening of Consumers with Non-common Priors,” mimeo, UCL.

Appendix

Proof of Proposition 1. Assume that the principal knows the value of θ . Then, an optimal contract is given by the solution to the problem,

$$\max_{t_\theta(a)} \{\theta [t_\theta(a^u) - c(a^u)] + (1 - \theta) [t_\theta(a^v) - c(a^v)]\}$$

subject to the IR_θ , UR_θ and VR_θ constraints. Because the principal does not need to worry about incentive compatibility constraints, IR_θ is binding. This means that the following contract maximizes the principal’s expected profits, while satisfying the IR_θ constraint.

$$t_\theta(a) = \begin{cases} u^* & \text{if } a = a_{\max}^{u-c} \\ v^* & \text{if } a = a_{\max}^{v-c} \\ \infty & \text{if } a \notin \{a_{\max}^{u-c}, a_{\max}^{v-c}\} \end{cases}$$

By the crossing property, this contract also satisfies the UR_θ and VR_θ constraints. Because this contract is independent of θ , it is also optimal when the principal does not know the value of θ and needs to take into consideration the incentive compatibility constraints. ■

Proof of Lemma 1. Assume the optimal menu includes a set of non-speculative contracts T (which may be a singleton) which generate an expected profit for the principal which is lower than his expected profit from the following contract t^* , which is a first-best contract against $\theta = p$:

$$t^*(a) = \begin{cases} u^* & \text{if } a = a_{\max}^{u-c} \\ v^* & \text{if } a = a_{\max}^{v-c} \\ \infty & \text{if } a \notin \{a_{\max}^{u-c}, a_{\max}^{v-c}\} \end{cases}$$

where a_{\max}^{u-c} and a_{\max}^{v-c} satisfy the condition in Definition 1. This contract has two important features: it yields zero expected payoff to all types, and it satisfies the UR and VR constraints of any type who chooses it. Let Θ^T denote the set of types whose most preferred contract in the menu is in T .

Consider amending the original menu by replacing all the contracts in T with t^* . Since the original menu is assumed to be optimal, every speculative contract in it must satisfy the UR and VR constraints. Hence, every contract in the new menu satisfies these constraints.

We wish to show that there is a way to assign a contract in the new menu to each agent type, such that (i) no agent type is assigned a contract with a negative expected payoff, (ii) the assigned contract is weakly preferred to any other contract in the menu, and (iii) the new menu would generate a higher expected profit than the original menu.

By our construction of t^* , this is straightforward if we can assign t^* to every type in Θ^T , while assigning all other types their most preferred contract in the original menu. Suppose there is a positive measure of types in Θ^T who strictly prefer a contract $t \notin T$ to the contract t^* . If t is non-speculative, then by the definition of T , the expected profit from t is larger or equal to the expected profit from t^* . If t is speculative, then its expected profit must be at least as high as the expected profit from t^* (otherwise, the principal would be able to increase his expected profits by replacing the original menu with one that consists of t^* and all the original speculative contracts except for the least profitable one).

Suppose next that there is a positive measure of types outside of Θ^T whose most preferred contract in the new menu is t^* . Then each of these types must have obtained a strictly negative expected payoff from his most preferred contract in the original menu, contradicting the *IR* constraint. ■

Proof of Lemma 2. Assume that $0 < p < 1$ and yet the optimal menu does *not* include any speculative contract. By the crossing condition, $u^* > v(a_{\max}^{u-c})$ or $v^* > u(a_{\max}^{v-c})$. Consider the former case. Choose ε such that $u^* - \varepsilon > v(a_{\max}^{u-c})$ and

$$0 < \frac{\varepsilon}{u^* - v(a_{\max}^{u-c}) - \varepsilon} < \frac{1-p}{p}$$

Then

$$p(u^* - \varepsilon) + (1-p)[v^* + u^* - v(a_{\max}^{u-c}) - \varepsilon] > pu^* + (1-p)v^* \quad (12)$$

In addition, there exists $\bar{\theta} < 1$ sufficiently close to one that satisfies

$$\bar{\theta}\varepsilon - (1-\bar{\theta})[u^* - v(a_{\max}^{u-c}) - \varepsilon] = 0 \quad (13)$$

Suppose that the principal added to the original menu the following speculative contract:

$$t(a) = \begin{cases} u^* - \varepsilon & \text{if } a = a_{\max}^{u-c} \\ v^* + u^* - v(a_{\max}^{u-c}) - \varepsilon & \text{if } a = a_{\max}^{v-c} \\ \infty & \text{if } a \notin \{a_{\max}^{u-c}, a_{\max}^{v-c}\} \end{cases}$$

By Lemma (1), the non-speculative contract in the optimal menu is t^* . By (13), all types $\theta > \bar{\theta}$ strictly prefer $t(a)$ to t^* , while the opposite is true of types below $\bar{\theta}$. By (12) and the continuity of F , the expected profit from the new menu is strictly higher than with the original menu, a contradiction.

By essentially the same argument, we can show that if $v^* > u(a_{\max}^{v-c})$, then we can design

a speculative contract with the following properties: (i) there exists $\underline{\theta} > 0$ such that all types $\theta < \underline{\theta}$ strictly prefer this contract to t^* , and (ii) the principal obtains an expected profit strictly above $p[u^* - c(a_{\max}^{u-c})] + (1-p)[v^* - c(a_{\max}^{v-c})]$. Taken together, these properties imply that the principal can raise his expected profit by including this speculative contract in his original menu, a contradiction.

Assume next that $p = 0$. Let us first prove the “if” part. Suppose that $u(a') > v(a')$ for some a' , yet the optimal menu has *no* speculative contracts. By the crossing property and the continuity of v and c , there exists $a_{\max}^{v-c} \neq a'$. The principal can then add the following contract to the original menu,

$$t(a) = \begin{cases} u(a') - \varepsilon & \text{if } a = a' \\ v^* + u(a') - v(a') - \varepsilon & \text{if } a = a_{\max}^{v-c} \\ \infty & \text{if } a \notin \{a', a_{\max}^{v-c}\} \end{cases}$$

where $0 < \varepsilon < u(a') - v(a')$. This contract yields the principal a higher profit than any non-speculative contract. In addition, there exists $\bar{\theta} < 1$ such that $t(a)$ yields a strictly positive expected payoff to all types $\theta > \bar{\theta}$ and a strictly negative expected payoff to all types below $\bar{\theta}$. By the continuity of F , the new menu generates a strictly higher expected profit than the original menu, a contradiction. Thus, if $p = 0$ and $u(a') > v(a')$ for some a' , then the optimal menu must contain at least one speculative contract.

Turning to the “only if” part, suppose that $u(a) \leq v(a)$ for every a . When the principal knows the agent’s type, his objective is to solve the following maximization problem:

$$\max_{t^u, t^v, a^u, a^v} t^v - c(a^v)$$

subject to

$$\theta[u(a^u) - t^u] + (1 - \theta)[v(a^v) - t^v] \geq 0 \quad (IR)$$

$$u(a^u) - t^u \geq u(a^v) - t^v \quad (UR)$$

$$v(a^v) - t^v \geq v(a^u) - t^u \quad (VR)$$

By assumption, $u(a) \leq v(a)$ for every a . The *VR* constraint thus implies $v(a^v) - t^v \geq u(a^u) - t^u$. By the *IR* constraint, $t^v \leq v(a^v)$. Therefore, the principal’s objective function is bounded from above by $v(a^v) - c(a^v)$. It follows that the non-speculative contract t^* maximizes the principal’s payoff, regardless of the agent’s type. Therefore, the optimal menu contains no speculative contracts.

An argument along the same lines establishes that when $p = 1$, the optimal menu includes at least one speculative contract if and only if, $u(a) < v(a)$ for some a . ■

Proof of Lemma 3. Note first that by *IR_p*, type $\theta = p$ cannot be assigned a speculative contract. Our proof relies on the following observation.

Observation 1. Assume the optimal menu assigns a speculative contract to a type $\theta \neq p$. Then $D_\theta^u > 0$ and $D_\theta^v < 0$ for $\theta > p$, while $D_\theta^u < 0$ and $D_\theta^v > 0$ for $\theta < p$.

Proof of Observation 1. Suppose $\theta > p$. Let $x \equiv \theta - p > 0$. Then by Definition 2,

$$(\theta - x) D_\theta^u + (1 - \theta + x) D_\theta^v < 0$$

By IR_θ ,

$$\theta D_\theta^u + (1 - \theta) D_\theta^v \geq 0$$

Because $\theta > x > 0$, these two inequalities imply that $D_\theta^u > 0$ and $D_\theta^v < 0$. A similar argument applies for $\theta < p$. \square

We now show that if a type $\theta > p$ is assigned a speculative contract, then a higher type must also be assigned such a contract. Consider a type $\phi > \theta$. Assume this type is assigned a non-speculative contract. Then by Lemma 1, $U(\phi, \phi) = 0$. By $IC_{\phi, \theta}$, an agent of type $\phi > \theta$ satisfies $U(\phi, \phi) \geq U(\theta, \phi)$. By Observation 1, $U(\cdot, x)$ is affine in x , hence, $U(\theta, \phi) > U(\theta, \theta)$. Since $U(\theta, \theta) \geq 0$, we reached a contradiction. A similar argument applies for types lower than p . \blacksquare

Proof of Corollary 1. Assume that $U(\bar{\theta}, \bar{\theta}) > 0$. By Observation 1, $U(\cdot, x)$ is affine in x . Hence, $U(\bar{\theta}, \phi) > 0$ for all $\phi \geq \bar{\theta}$. Suppose the principal deviates by modifying all the speculative contracts as follows: for every $\phi \geq \bar{\theta}$, $t(a_\phi^u)$ and $t(a_\phi^v)$ are both raised by some arbitrarily small ε . This modification leaves all the IR , IC , UR and VR constraints intact, and generates a higher profit, a contradiction. \blacksquare

Proof of Proposition 2. Assume $u(a) \geq v(a)$ for all a . By Lemma 2, if $p = 1$, there are no speculative contracts in the optimal menu. Now let $p < 1$. Assume, by contradiction, that type $\theta = 0$ is assigned a speculative contract t_0 . Then there exists a pair of actions (a, a') such that $v(a) \geq t_0(a)$ and $u(a') < t_0(a')$. But from our assumption that $u(a) \geq v(a)$ it follows that $u(a') - t_0(a') < u(a) - t_0(a)$, which violates UR_0 , a contradiction. By Lemma 3, if $\theta = 0$ is not assigned a speculative contract, then no type $0 < \theta < p$ can be assigned such a contract. A similar proof establishes the second part of the result. \blacksquare

Proof of Lemma 4. Assume that the principal designs two separate menus such that one is optimal for the distribution of types F conditional on $\theta \in [p, 1]$, and another is optimal for the distribution F conditional on $\theta \in [0, p]$. Denote the first set of contracts by T^+ and the second by T^- . We claim that these menus have the property that each type in $[p, 1]$ (respectively $[0, p]$) weakly prefers his assigned contract to every contract in T^- (respectively T^+).

By Lemmas 1 and 3, there exist $\underline{\theta} \in [p, 1]$ and $\bar{\theta} \in [0, p]$ such that $U(\theta, \theta) = 0$ for all $\theta \in (\underline{\theta}, \bar{\theta})$. By Observation 1, $D_\theta^v > D_\theta^u$ for all $\theta \leq \underline{\theta}$, while $D_\theta^u > D_\theta^v$ for all $\theta \geq \bar{\theta}$. In addition,

every $\theta \notin (\underline{\theta}, \bar{\theta})$ satisfies inequality (2). This means that for every pair of types θ, ϕ such that $\theta \leq p$ and $\phi \geq \bar{\theta}$ we have that $U(\phi, \theta) < 0$. Similarly, $U(\phi', \theta') < 0$ for every pair of types θ', ϕ' with $\theta' \geq p$ and $\phi \leq \underline{\theta}$.

By assumption, T^+ (respectively T^-) maximizes the principal's expected profit, conditional on $\theta \in [p, 1]$ (respectively $\theta \in [0, 1]$). Hence, the union of T^+ and T^- maximizes the *unconditional* expected profit of the principal. In addition, T^+ satisfies the *IR*, *IC*, *UR* and *VR* constraints of the types in $[p, 1]$, while T^- satisfies the corresponding constraints for types in $[0, p]$. From the argument made in the previous paragraph, the set of contracts $T^+ \cup T^-$ also satisfies these constraints for *all* types in $[0, 1]$. ■

Proof of Lemma 5. Consider the types in $[p, 1]$. By Lemma 3, there exists $\bar{\theta} \in [p, 1]$ such that all types $\theta < \bar{\theta}$ choose a non-speculative contract. Hence, by Lemma 1, $a_\theta^v \in \arg \max_a [u(a) - c(a)]$.

By Lemma 3, if $\bar{\theta} < 1$, then all types $\theta \geq \bar{\theta}$ are assigned speculative contracts. Assume there exists $\theta^* > \bar{\theta}$ who is assigned a contract with $a_\theta^v \notin \arg \max_a [v(a) - c(a)]$. Consider the modified contract t'_θ , which differs from t_θ only in two actions, a_θ^v and some a_{\max}^{v-c} (as defined at the end of Section 2), such that $t'(a_\theta^v) = \infty$ and $t'_\theta(a_{\max}^{v-c})$ satisfies

$$v(a_{\max}^{v-c}) - t'_\theta(a_{\max}^{v-c}) = v(a_\theta^v) - t_\theta(a_\theta^v) \quad (14)$$

Because the agent's net payoff in state v from the modified contract is the same as under the original contract, all the *IR*, *IC* constraints, as well as *VR* $_\theta$, continue to hold. To see that *UR* $_\theta$ is also satisfied, note that because t_θ is speculative and $\theta > p$, $D_\theta^u > 0$ and $D_\theta^v < 0$, by Observation 1. Therefore, by (14), $u(a_\theta^v) - t(a_\theta^v) \geq v(a_{\max}^{v-c}) - t'_\theta(a_{\max}^{v-c})$.

Finally, conditional on the agent's type being θ , the contract t'_θ generates a higher expected profit to the principal. This follows from observing that (14), together with our assumption on a_θ^v , implies that $t'_\theta(a_{\max}^{v-c}) - c(a_{\max}^{v-c}) > t_\theta(a_\theta^v) - c(a_\theta^v)$.

A similar argument shows that we may set $a_\theta^u \in \arg \max_a [u(a) - c(a)]$ for $\theta \leq p$. ■

Proof of Lemma 7. We adopt Krishna's (2002, pp. 63-66) derivation of incentive compatibility for direct mechanisms. Define $m(\theta) \equiv -D_\theta^v$. The optimal menu is incentive compatible if for all types θ and ϕ ,

$$V(\theta) \equiv \theta q(\theta) - m(\theta) \geq \phi q(\theta) - m(\phi)$$

By Observation 1, $q(\theta) \geq 0$ for all $\theta \geq \underline{\theta}$ (by the definition of $\underline{\theta}$, $q(\theta) = 0$ for all $\theta < \underline{\theta}$). Hence, the L.H.S of the above inequality is an affine function of the true value θ . Incentive compatibility implies that for all $\theta \geq \bar{\theta}$,

$$V(\theta) = \max_{\phi \in [0, 1]} \{\theta q(\phi) - m(\phi)\}$$

I.e., $V(\theta)$ is a maximum of a family of affine functions, and hence it is convex on $[\bar{\theta}, 1]$.⁵

Incentive compatibility is equivalent to the requirement that for all $\theta, \phi \in [\bar{\theta}, 1]$,

$$V(\phi) \geq V(\theta) + q(\theta)(\phi - \theta)$$

This implies that for all $\theta > \bar{\theta}$, $q(\theta)$ is the slope of a line that supports the function $V(\theta)$ at the point θ . Because $V(\theta)$ is convex it is absolutely continuous, and thus differentiable almost everywhere in the interior of its domain. Hence, at every point that $V(\theta)$ is differentiable, $V'(\theta) = q(\theta)$. Since $V(\theta)$ is absolutely continuous, we obtain that for all $\theta > \bar{\theta}$,

$$V(\theta) = V(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(x) dx$$

By Corollary 1, $V(\underline{\theta}) = 0$, and so we obtain that $U(\theta, \theta) = \int_{\underline{\theta}}^{\theta} q(x) dx$ for $\theta \geq \bar{\theta}$.

By a similar argument, if the speculative contract assigned to type $\theta \leq \underline{\theta}$ is speculative, then $U(\theta, \theta) = \int_{1-\underline{\theta}}^{1-\theta} [-q(x)] dx$. ■

Proof of Lemma 8. Consider first the types $\theta \geq \bar{\theta}$. Let δ_{θ} denote the slack in the VR_{θ} constraint of type θ . Then

$$t_{\theta}(a_{\theta}^v) - t_{\theta}(a_{\theta}^u) = v(a_{\theta}^v) - v(a_{\theta}^u) - \delta_{\theta} \quad (15)$$

By Lemma 7, if type $\theta \geq \bar{\theta}$ is assigned a speculative contract, which is incentive compatible, then

$$U(\theta, \theta) = \int_{\bar{\theta}}^{\theta} [u(a_x^u) - v(a_x^u) - \delta_x] dx \quad (16)$$

Assume the optimal menu satisfies that $\delta_{\theta} > 0$ for some positive measure of types in $[\bar{\theta}, 1]$. Consider amending the menu by changing only $t_{\theta}(a_{\theta}^v)$ for all types $\theta \geq \underline{\theta}$ such that the new transfer is equal to $t_{\theta}(a_{\theta}^v) + \delta_{\theta}$, making the VR_{θ} constraint binding for all these types. Clearly, this change does not violate the UR_{θ} constraint of these types. From (16), it follows that the incentive compatibility constraints are not violated, and that this change only raises $U(\theta, \theta)$. Hence, the IR_{θ} constraint is also not violated.

We claim that the above change in the menu, increases the principal's expected profit, in contradiction to our assumption that the original menu was optimal. To see this, note that by Lemma 3, the optimal menu maximizes the expression

$$\begin{aligned} & \left[\frac{F(\bar{\theta}) - F(p)}{1 - F(p)} \right] [p(u^* - c(a_{\max}^{u-c})) + (1-p)(v^* - c(a_{\max}^{v-c}))] \\ & + \int_{\bar{\theta}}^1 [(1-p)(t_{\theta}^v - t_{\theta}^u - c_{\theta}^v + c_{\theta}^u) + t_{\theta}^u - c_{\theta}^u] \left[d \frac{F(\theta)}{1 - F(p)} \right] \end{aligned}$$

⁵Because all types in $[p, \bar{\theta}]$ are assigned a non-speculative contract, $V(\theta) = 0$ for all $\theta \in [p, \bar{\theta}]$.

Substituting (15) into the above expression, we obtain that reducing δ_θ for a positive measure of types in $[\bar{\theta}, 1]$ increases the principal's expected profit.

A similar argument establishes that the UR_θ constraint is binding for all types $\theta \leq \underline{\theta}$. ■

Proof of Proposition 3. Assume that F satisfies the monotone hazard rate condition. By Lemma 5, we may restrict attention to menus that induce $a_\theta^v \in \arg \max_a [v(a) - c(a)]$ for $\theta \geq p$, and $a_\theta^u \in \arg \max_a [u(a) - c(a)]$ for $\theta \leq p$. Consider first the problem of designing an optimal menu for types distributed on $[p, 1]$. By the lemmas we have proven thus far, we may reduce this problem to the following optimization problem:

$$\begin{aligned} & \max_{\bar{\theta}, \{a_\theta^u, t_\theta\}_{\theta \in [\bar{\theta}, 1]}} \left[\frac{F(\bar{\theta}) - F(p)}{1 - F(p)} \right] [p(u^* - c(a_{\max}^{u-c})) + (1-p)(v^* - c(a_{\max}^{v-c}))] \\ & + \int_{\bar{\theta}}^1 \{(1-p)[v^* - c(a_{\max}^{v-c}) - v(a_\theta^v) + c(a_\theta^u)] + t_\theta(a_\theta^u) - c(a_\theta^u)\} \left[d \frac{F(\theta)}{1 - F(p)} \right] \end{aligned}$$

subject to the IR_θ , $IC_{\theta, \phi}$, UR_θ and VR_θ constraints for all types θ and ϕ . We adopt the standard practice in mechanism-design (see Krishna (2002)) of solving this problem under the assumption that the $IC_{\theta, \phi}$ and IR_θ constraints are satisfied, and then checking that this constraint is indeed satisfied by the solution we obtain.

Because VR_θ is binding (Lemma 8) for all $\theta \geq \bar{\theta}$, we can write

$$t_\theta(a_\theta^u) = t_\theta(a_\theta^v) + v(a_\theta^u) - v^* \quad (17)$$

Equating (3) and (4), we can obtain the following expression for $t_\theta(a_\theta^v)$,

$$t_\theta(a_\theta^v) = \theta q(\theta) + v^* - \int_{\bar{\theta}}^{\theta} q(x) dx \quad (18)$$

Substituting this expression into (17), we obtain

$$t_\theta(a_\theta^u) = \theta q(\theta) - \int_{\bar{\theta}}^{\theta} q(x) dx + v(a_\theta^u) \quad (19)$$

We can thus simplify the optimization problem by expressing the objective as a function only of the cutoff $\bar{\theta}$ and the actions $\{a_\theta^u\}_{\theta \in [\bar{\theta}, 1]}$,

$$\max_{\bar{\theta}, \{a_\theta^u\}_{\theta \in [\bar{\theta}, 1]}} (F(\bar{\theta}) - F(p))p(u^* - c(a_{\max}^{u-c})) + \int_{\bar{\theta}}^1 \left\{ \theta q(\theta) - \int_{\bar{\theta}}^{\theta} q(x) dx + p[v(a_\theta^u) - c(a_\theta^u)] \right\} dF(\theta)$$

Note that we have taken the constant $(1-p)(v^* - c(a_{\max}^{v-c}))$ out of the objective function. By interchanging the order of integration in $\int_{\bar{\theta}}^1 \int_{\bar{\theta}}^{\theta} q(x) dx$, we can rewrite this expression as $\int_{\bar{\theta}}^1 [1 - F(\theta)] q(\theta) d\theta$. We can thus replace $\int_{\bar{\theta}}^1 \left[\theta q(\theta) - \int_{\bar{\theta}}^{\theta} q(x) dx \right] dF(\theta)$ with $\int_{\bar{\theta}}^1 \psi_F(\theta) q(\theta) dF(\theta)$, where $\psi_F(\theta)$ is defined in (5). Because VR_θ is binding we can replace

$q(\theta)$ with $u(a_\theta^u) - v(a_\theta^u)$ to obtain the following objective function

$$\max_{\bar{\theta}, \{a_\theta^u\}_{\theta \in [\bar{\theta}, 1]}} (F(\bar{\theta}) - F(p))p(u^* - c(a_{\max}^{u-c})) + \int_{\bar{\theta}}^1 \{\psi_F(\theta)(u(a_\theta^u) - v(a_\theta^u)) + p[v(a_\theta^u) - c(a_\theta^u)]\} dF(\theta) \quad (20)$$

Because F satisfies the monotone hazard rate property, it follows that the optimal a_θ^u for each $\theta \in [\bar{\theta}, 1]$ is given by (6). A similar argument establishes that the optimal a_θ^v for each $\theta \in [0, \bar{\theta}]$ is given by (7).⁶

It remains to verify that the menu we constructed is indeed individually-rational and incentive compatible. Note first that the non-speculative contract t^* gives zero indirect utility to types in $(\underline{\theta}, \bar{\theta})$, while speculative contracts, by definition, give negative indirect utility to these types. We now show that the speculative contracts assigned to types above p satisfy the *IR* and *IC* constraints of these types. Analogous arguments apply to speculative types below p .

By Corollary 1, the *IR* constraint of the lowest speculative type is binding. Hence, the speculative contracts assigned to all higher types is individually rational if, and only if they are incentive compatible. Incentive compatibility is equivalent to the requirement that $q(\theta)$ is non-decreasing (see Krishna (2002), p.68). Since by Lemma 8, VR_θ is binding for all $\theta \geq \bar{\theta}$, $q(\theta)$ is non-decreasing for these types if, and only if $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing for all $\theta \geq \bar{\theta}$. Consider a pair of types ϕ, θ such that $\phi > \theta$. By construction,

$$\psi_F(\phi)[u(a_\phi^u) - v(a_\phi^u)] + p[v(a_\phi^u) - c(a_\phi^u)] \geq \psi_F(\phi)[u(a_\theta^u) - v(a_\theta^u)] + p[v(a_\theta^u) - c(a_\theta^u)]$$

and

$$\psi_F(\theta)[u(a_\theta^u) - v(a_\theta^u)] + p[v(a_\theta^u) - c(a_\theta^u)] \geq \psi_F(\theta)[u(a_\phi^u) - v(a_\phi^u)] + p[v(a_\phi^u) - c(a_\phi^u)]$$

Adding these two inequalities and cancelling common terms yields

$$[\psi_F(\phi) - \psi_F(\theta)] \{[u(a_\phi^u) - v(a_\phi^u)] - [u(a_\theta^u) - v(a_\theta^u)]\} \geq 0$$

Because F satisfies the monotone hazard rate property, $\psi_F(\phi) > \psi_F(\theta)$. Therefore, $u(a_\phi^u) - v(a_\phi^u) \geq u(a_\theta^u) - v(a_\theta^u)$. ■

Proof of Proposition 4. We begin by proving part (i) of the proposition. Define

$$h(\theta) \equiv \max_{a \in [0, 1]} [(\psi_F(\theta) - p) \cdot (u(a) - v(a)) + p \cdot (u(a) - c(a))] - p \cdot (u^* - c(a_{\max}^{u-c})) \quad (21)$$

and let

$$a_\theta \in \arg \max_{a \in [0, 1]} [(\psi_F(\theta) - p) \cdot (u(a) - v(a)) + p \cdot (u(a) - c(a))]$$

⁶Note that in this case we need to use $G(z)$ instead of $F(\theta)$.

It follows from the proof of Proposition 3 that the principal strictly prefers to speculate with agent type θ if and only if $h(\theta) > 0$. By the monotone hazard rate property, $\psi_F(\theta)$ is a continuously increasing function. Note that $\psi_F(p) < p$ and $\psi(1) = 1$. Therefore, there exists a unique $\theta^* > p$ that solves the equation $\psi_F(\theta) = p$. By Lemma 6, $q(\theta) > 0$ for $\theta \geq \bar{\theta}$. By Lemma 8, VR is binding for $\theta \geq \bar{\theta}$. By the definition of $q(\theta)$, this means that $u(a_\theta) - v(a_\theta) > 0$ for $\theta \geq \bar{\theta}$.

Consider $\theta \geq \bar{\theta}$ for which $\psi_F(\theta) < p$. Because $u(a) - c(a) < u^* - c(a_{\max}^{u-c})$, the only way that $h(\theta) > 0$ is if $u(a) - v(a) < 0$, a contradiction. Now let $\theta = \theta^*$. Then $h(\theta^*) = 0$. But since $\psi_F(\theta)$ is increasing with θ , this means that $h(\theta) > 0$ for every $\theta > \theta^*$. Therefore, $\bar{\theta} = \theta^*$.

A similar argument establishes part (ii) of the proposition. ■

Proof of Proposition 5. Consider first the speculative contracts in the optimal menu for types higher than p . We already know from Lemma 5 that w.l.o.g. we can set $a_\theta^v = a_{\max}^{v-c}$ for all $\theta \geq \bar{\theta}$. By ex-post efficiency, we can let $a_\theta^u = a_{\max}^{u-c}$ for all $\theta \geq \bar{\theta}$. By Lemma 6, this means that $q(\theta) = u^* - v(a_{\max}^{u-c})$ for all types $\theta \geq \bar{\theta}$. Hence, from equations (17) and (18) it follows that the transfers $t_\theta(a_\theta^u)$ and $t_\theta(a_\theta^v)$ are also constant for all $\theta \geq \bar{\theta}$. By an analogous argument it can be shown that $t_\theta(a_\theta^u)$ and $t_\theta(a_\theta^v)$ are constant for all $\theta \leq \underline{\theta}$. ■

Proof of Proposition 6. Suppose $a_\theta^u \notin \arg \max_a [u(a) - c(a)]$. Then by assumption, $a_\theta^u \notin \arg \max_a [u(a) - v(a)]$. Consider a modified contract t'_θ , which differs from t_θ in two actions only, a_θ^u and some a_{\max}^{u-c} , such that $t'_\theta(a_\theta^u) = \infty$, and $t'_\theta(a_{\max}^{u-c}) = t_\theta(a_\theta^u) + u(a_{\max}^{u-c}) - u(a_\theta^u)$. Note that $t'_\theta(a_{\max}^{u-c}) - c(a_{\max}^{u-c}) > t_\theta(a_\theta^u) - c(a_\theta^u)$, hence, conditional on being chosen, this contract generates a higher expected profit than t_θ . By construction, the agent's net payoff in state u is the same under t_θ and t'_θ . Therefore, all the IR and IC constraints, as well as UR_θ , continue to hold. Since, by assumption, $a_{\max}^{u-c} \in \arg \max_a [u(a) - v(a)]$, and since the original menu satisfied the VR_θ constraint, it can easily be checked that this constraint continues to hold. But this contradicts our assumption that the original menu was optimal. A similar argument shows that $a_\theta^v \in \arg \max_a [v(a) - c(a)]$. ■