

Coordination under Limited Depth of Reasoning

Terri Kneeland*

November 4th, 2015

Abstract

This paper provides the first unified explanation of behavior in coordinated attack games under both public and private information. It demonstrates that the main experimental results, such as threshold strategies, comparative statics, and the differences in behavior under public and private information, are robust predictions of limited depth of reasoning models. This is in contrast to equilibrium, which mispredicts the coordinating roles of public and private information. The analysis has implications for understanding macroeconomic phenomena, like currency attacks and debt crises, which are commonly modeled using incomplete information coordinated attack games.

Keywords: coordination games, level-k models, cognitive hierarchy models, global games

1 Introduction

Consider a simple coordination game used to model a speculative attack. Players have the option to attack a currency peg. If enough players attack, the attack is successful, and the peg collapses. If not enough players attack, the attack fails, and the peg holds. The threshold for a successful attack depends on the fundamentals of

*Department of Economics, University College London, Gower Street, London, UK WC1E 6BT (e-mail: t.kneeland@ucl.ac.uk). I would like to thank my advisor, Yoram Halevy, for substantial guidance and feedback. I am grateful to Colin Camerer, Vince Crawford, David Freeman, Li Hao, Frank Hienemann, Wei Li, Mike Peters, Luba Petersen, and Matthew Rabin for helpful comments and discussions. And, I am grateful to Frank Heinemann, Rosemarie Nagel and Peter Ockenfels for allowing me access to their experimental data. Comments by anonymous referees were greatly appreciated.

the economy. If fundamentals are strong, a large proportion of players must attack in order for the attack to be successful. If fundamentals are weak, only a small proportion of players are needed for a successful attack.

Heinemann et al. (2004), Duffy and Ochs (2012), Heinemann et al. (2009), and Cornand (2006) test coordinated attack games in the laboratory. There are two main aggregate findings. (1) Players use threshold strategies. They attack when fundamentals are weak and restrain from attacking when fundamentals are strong. This holds regardless of whether players receive public or private information about fundamentals. And, (2) the degree of coordination depends upon the information structure.¹ Public information increases coordination relative to private information.

These results provide a challenge for equilibrium theory. Private information generates disperse higher-order beliefs which weaken complementarities in actions generating a unique threshold equilibrium. In contrast, public information strengthens complementarities between actions generating multiple equilibria driven by self-fulfilling beliefs. Equilibrium predicts the use of threshold strategies under private information but does not explain why behavior is tied to fundamentals under public information. Further, equilibrium mispredicts the coordinating roles of different types of information, predicting weakly better coordination under private information than under public information.²

This paper shows that an alternative solution concept, based on limited depth of reasoning, provides a unified explanation of the pattern of behavior under both public and private information. To the best of our knowledge, this is the first paper to provide a consistent explanation of the existing experimental literature on coordinated attack games.

Limited depth of reasoning, like level-k thinking and cognitive hierarchy, is a behaviorally-motivated approach to reasoning in games.³ Each player has a bounded

¹We use the term coordination here to mean that players use the same strategies. In a complete information game, if players use the same pure strategy they would perfectly coordinate on either attacking or not attacking. But under incomplete information, even if players use the same strategy, there could be ex post coordination failure where some players attack and others do not attack because they have different private information. However, if players are not using the same threshold strategy this will only increase the degree of ex post coordination failure.

²Even though equilibrium allows multiple equilibria under public information, it is possible for players to coordinate on the same equilibrium, thus generating perfect coordination. The inconsistency of equilibrium with experimental data lies in the private information predictions. Equilibrium fails to explain the observed coordination failure under private information.

³For, pioneering works in the literature see Stahl and Wilson (1994; 1995), Nagel (1995), Costa-

depth of reasoning determined by her cognitive type. L0 types have zero depths of reasoning - their behavior is specified outside of the model. L1 types have one depth of reasoning. They best respond to L0 types. L2 types have two depths of reasoning. They best respond to some belief over L0 and L1 types. And, so on, with Lk types best responding to some belief over lower types $\{L0, L1, \dots, Lk-1\}$. All higher types anchor their beliefs in the nonstrategic L0 type and calculate their optimal actions using a finite number of iterated best responses.

Limited depth of reasoning is consistent with the existing experimental results. Behavior is tied to fundamentals under any information structure with players attacking only when fundamentals are weak. Moreover, limited depth of reasoning predicts that coordination is greater under public information than private information. Players with different depths of reasoning treat public information similarly but private information differently. Public information coordinates higher- and lower-order beliefs about fundamentals. This coordinates the behavior of players with different depths of reasoning. In contrast, private information decreases coordination by creating differences between higher- and lower-order beliefs and hence differences in the behavior of players with different depths of reasoning.

This paper goes on to show that the aggregate experimental results are robust predictions of limited depth of reasoning models. The results hold under weak assumptions on the non-strategic L0 type. And, they hold under any limited depth of reasoning model. That is, they hold for any specification of beliefs that an Lk type might hold over lower types and for any distribution of types in the model. The results even hold if there is some proportion of sophisticated types in the model (types that have correct beliefs about the types of others, infinite depths of reasoning, and realize that others may also have infinite depths of reasoning).

Part of the difficulty with models like level-k thinking and cognitive hierarchy is that L0 behavior and the type distribution must be specified outside of the model. This is problematic because both L0 behavior and the type distribution do not appear to be stable across games. However, the theoretical predictions of this paper are largely independent of these specifications. Even without knowing L0 behavior or the type distribution, limited depth of reasoning models make robust, testable predictions.

In addition, this paper analyzes the experimental data from Heinemann et al.

Gomes et al. (2001), and Camerer et al. (2004). For a recent survey of this literature, see Costa-Gomes et al. (2013).

(2004) and classifies subjects as either level-k or equilibrium types using maximum likelihood to fit a finite mixture model to the data. Allowing for level-k types significantly improves the model’s ability to explain the experimental data. Approximately, 83 percent of subjects can be classified as level-k types and 17 percent of subjects as equilibrium types. Further, the level-k model explains the data just as well as a model with level-k and equilibrium types. Limited depth of reasoning provides a consistent explanation for experimental behavior at both the individual and aggregate level.

Three papers study a related class of incomplete information games under limited depth of reasoning. Shapiro et al. (2014) and Cornand and Heinemann (2014) test these games in the laboratory and find, respectively, that the level-k and cognitive hierarchy models are consistent with behavior.⁴ The analysis is similar to the maximum likelihood analysis done in this paper, showing that there exists a limited depth of reasoning model that fits the experimental behavior better than equilibrium. We reinforce these results by demonstrating that the same holds true for the coordinated attack game. But, the insights from these papers can not be extended to the coordinated attack game. Neither of the two papers investigates the coordinating roles of different types of information or considers the robust predictions of limited depth of reasoning models. Cornand and Heinemann (2015) uses the estimated type distribution from Cornand and Heinemann (2014) to analyze the welfare effects of public information under the cognitive hierarchy model.

The motivation in this paper is similar to Strzalecki (2014) who uses limited depth of reasoning models to explain observed behavior in Rubinstein’s (1989) email game, which is a 2-player coordinated attack game with private information. Strzalecki investigates whether the limited depth of reasoning model is consistent with observed behavior under a particular information structure. This paper is focused on understanding behavior across different information structures. In addition, the limited depth of reasoning solution concept is applied differently in the two papers. The analysis of Strzalecki would be equivalent to fixing the solution concept as the set m^{th} -order rationalizable actions (for some fixed and small m). However, the analysis in the current paper closely follows the approach of the existing level-k and cognitive hierarchy literature and focuses on the outcomes that result from a fixed $L0$.

⁴Players have both an information motive (incentive to match fundamentals) and a coordination motive (incentive to match behavior of others) in these games. Shapiro et al. find that the level-k model does particularly well in games where the coordination motive dominates the information motive, and less well when the information motive dominates.

This paper is closely related to the global games literature. The coordinated attack game analyzed here is one of the workhorse models of global games. Global games, initiated by Carlsson and van Damme (1993b,a) and furthered by Morris and Shin (1998, 2003), has shown that multiplicity in coordination games stems from common knowledge of payoffs. The concern is that increased precision of public information may be destabilizing; it could lead to multiplicity driven by self-fulfilling beliefs. However, both the existing experimental literature and the boundedly rational approach to strategic reasoning studied here suggests that public information does not have destabilizing effects.⁵

This paper proceeds as follows. The next section describes the coordinated attack game and the equilibrium results. Section 3 establishes the limited depth of reasoning results. Section 4 analyzes the experimental data and discusses the consistency of level-k with the experimental literature. And, Section 5 concludes. Omitted proofs can be found in Appendix A.

2 Coordinated Attack Game

A simple model of a coordinated attack is presented. The model follows Morris and Shin (2004) and Bannier (2002). The coordinated attack game is interpreted as a speculative attack throughout the rest of this paper. This model can be applied to describe other phenomena such as bank runs, debt crises, coordinated investment, and political change.

There is a continuum of players indexed by i and uniformly distributed on $[0, 1]$. Players may attack the exchange rate peg or do nothing. There is a cost t of attacking. If a player attacks and the peg is abandoned, the player receives a positive payoff D , for a net payoff of $D - t$. It is assumed $D > t$. The payoff from not attacking is zero. The exchange rate peg is abandoned if and only if the proportion of players attacking, denoted by l , is no less than a critical value $\theta \in \mathbb{R}$.

The critical value θ parametrizes the strength of the status quo. It is often referred

⁵A second application of global games is as an equilibrium selection mechanism. In this application, all commonly known information is treated as private information (with vanishing noise). This selects a unique threshold equilibrium - providing an explanation for the observed behavior. However, this interpretation necessarily limits the study of behavior under different types of information because public information is equated with private information. Allowing for the existence of boundedly rational types allows us to talk meaningfully about behavior under different types of information.

to as the fundamentals of the economy. Under the speculative attack interpretation, θ can be interpreted as foreign exchange reserves held by the Central Bank. A higher θ represents better fundamentals and raises the threshold for a successful attack. The payoffs for each player can be summarized by the function $\pi : [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}$, the payoff gain from attacking. It is defined as

$$\pi(l, \theta) = \begin{cases} D - t & \text{if } l \geq \theta \\ -t & \text{otherwise} \end{cases}$$

2.1 Common Knowledge of Payoffs

There are three different cases to consider when payoffs are common knowledge:

$\theta > 1$ Even if all players attack, the fundamentals are sufficiently strong to maintain the exchange rate peg. There is a unique equilibrium in which none of the players attack.

$\theta \leq 0$ The fundamentals are too weak for the peg to be maintained. The unique equilibrium is one in which all players attack.

$0 < \theta \leq 1$ The currency regime is ripe for attack. There exist two equilibria - one in which all players attack and the exchange rate peg is abandoned and another in which no player attacks and the peg is upheld.

This game has multiple equilibria whenever $0 < \theta \leq 1$. Attacking is only worthwhile if a player expects others to attack. As beliefs are self-fulfilling, they are not tied to fundamentals and there is no way to predict an attack.

2.2 Incomplete Information of Payoffs

The above game can be converted into an incomplete information game by letting θ be unknown and having each player receive signals about θ . The distribution of θ and the signal processes are assumed to be common knowledge. Each player receives a private signal x_i and a public signal y about θ , where $x_i = \theta + \epsilon_i$ with $\epsilon_i \sim N(0, \frac{1}{\beta})$ and $y = \theta + \eta$ with $\eta \sim N(0, \frac{1}{\alpha})$. The fundamental θ is distributed uniformly on the

real line.^{6,7}

Conditional on the information received, θ is normally distributed with a mean formed by a weighted average of the public and private signals. Given signals x and y , θ is conditionally distributed according to

$$\theta|x, y \sim N\left((1 - \mu)y + \mu x, \frac{1}{\alpha + \beta}\right) \quad (1)$$

where $\mu = \frac{\beta}{\alpha + \beta}$. A player's strategy in the incomplete information game is a function $s : \mathbb{R} \rightarrow \{\text{attack, not attack}\}$, which lists an action for any private signal she might receive.

There is a unique equilibrium in the incomplete information game provided private information is sufficiently precise relative to public information (i.e. $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$). In this case, there will be a unique threshold equilibrium where players attack if and only if $x \leq \bar{x}^E$, for some unique \bar{x}^E .⁸ However, if public information is relatively precise there exists multiple equilibria driven by self-fulfilling beliefs. These results are established in Morris and Shin (2004).

If there is only private information (public information is imprecise, $\alpha \rightarrow 0$), then the condition, $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$, holds automatically and there is a unique threshold equilibrium. If there is only public information (private information is imprecise, $\beta \rightarrow 0$), the condition does not hold and there are multiple equilibria. Thus, equilibrium theory makes a sharp prediction for the case of private information but is uninformative about behavior under public information. Further, equilibrium mispredicts the coordinating roles of different types of information, predicting perfect coordination under private information.⁹

⁶The assumption that θ is uniformly distributed on the real line is typical of the global games literature. The assumption presents no technical difficulties as long as we restrict attention to conditional beliefs. See Hartigan (1983) for further discussion of this issue.

⁷The assumptions of normality and the uniform prior are not essential. The results go through with a generalized prior and signal structure. This is shown in a previous version of this paper, Kneeland (2012).

⁸Notice that the threshold cutoff $\bar{x}^E(y)$ will depend on the public signal y . We suppress that dependence and just write \bar{x}^E , taking the public signal y as fixed.

⁹The global games literature establishes that a unique equilibrium in these types of coordinated attack games depends on the degree of payoff uncertainty in the model (e.g. Morris and Shin, 2003). Typically, this is modeled by introducing asymmetric information over fundamentals. However, we can think about introducing payoff uncertainty in other ways, such as uncertainty about the risk preferences of other players. (Hellwig, 2002) shows that allowing for asymmetric information about the degree of risk aversion will induce a unique threshold equilibrium even when information about fundamentals is public information. Thus, this type of payoff uncertainty could explain why subjects

In the next section, we analyze the coordinated attack game under limited depth of reasoning and show that it's consistent with the existing experimental results.

3 Level-k Thinking

In a limited depth of reasoning model each player's behavior is determined by her cognitive type, which is drawn from a discrete distribution over a particular hierarchy of types $\{L0, L1, \dots, Lk, \dots\}$. L0 types are non-strategic, their behavior is specified outside the model. L0 types act as the starting point for players' strategic thinking: L1 best responds to L0, L2 to a distribution over L0 and L1, and so on. The behavior of all types can be calculated recursively, in a finite number of steps.

In a general limited depth of reasoning model, Lk types best respond to some distribution over lower types, where each type's beliefs about others are defined according to some prespecified rule. Under cognitive hierarchy, Lk types best respond to a mixture of L0, L1, \dots , Lk - 1 types with weights determined by the conditional Poisson distribution (e.g. Camerer et al. 2004). Under level-k thinking, Lk types best respond to Lk-1 types (e.g. Costa-Gomes and Crawford 2006). The main portion of this paper works within a level-k thinking model. This is for simplicity only. The results hold for any limited depth of reasoning model with Lk types best responding to some distribution over lower types (see Appendix B).

Both limited depth of reasoning and equilibrium require that players play a best response given their beliefs about others. This requirement alone is not enough to generate precise predictions. Equilibrium adds the additional assumption that players' actions and beliefs must be mutually consistent. Limited depth of reasoning models do not impose mutual consistency. Instead, they impose the assumption that players follow decision rules based on an iterated process of strategic thinking; players recursively calculate optimal behavior based on the anchoring L0 type. This defines a procedural model of player's decisions that avoids the circular logic of equilibrium imposed by the assumption of mutual consistency.

The estimated distribution of cognitive types tends to put most of the weight on L1 and L2 types and negligible weight on L0 types (Costa-Gomes and Crawford

coordinate successfully under public information. However, this type of explanation does not explain the failure to coordinate when information is private. Introducing other forms of payoff uncertainty will not increase coordination under public information while decreasing coordination under private information.

2006; Costa-Gomes et al. 2001). The anchoring L0 type exists mainly in the minds of others. This paper takes this position and assumes there is no support on L0 in the type distribution. In most applications, the specification of L0 is the key to the model’s explanatory power. However, the predictive power in this paper comes from the recursive nature of the level-k model. The main results hold under only weak restrictions on L0 behavior.

Let the behavior of L0 types be described by the cumulative distribution function $Q(l|x, y)$ on $[0, 1]$ and its associated density function $q(l|x, y)$. The distribution Q represents beliefs about the proportion of L0 types attacking. An L1 type who receives private information x and public information y believes that others are behaving according to $Q(l|x, y)$. The perceived behavior of L0 may be influenced by information.¹⁰ This is reasonable, as L0 types exist mainly in the minds of others. Two restrictions are placed on the behavior of L0.

A1 $Q(l|x, y)$ is weakly increasing in x and y for a given l .

A2 $q(l|x, y)$ is continuous in x and y for a given l .

Assumption A1 assumes that the perceived behavior of L0 types varies monotonically as information varies. This captures the likelihood that an L1 type with a low signal believes that others are more likely to attack than an L1 type with a higher signal. This is the natural specification for L0 as L0 is meant to capture players’ intuitive responses to the game. Assumption A2 requires $q(l|x, y)$ to be continuous with respect to x and y . This restriction is included for convenience to ensure continuity in the payoff functions. In previous level-k applications with incomplete information, L0 behavior was specified to be independent of information (e.g. Brocas et al. 2014 and Crawford and Iriberri 2007). Assumptions A1 and A2 are satisfied automatically in this case.

An L1 type who observes signals x and y believes that the behavior of others can be described by the density function $q(l|x, y)$. Given this, an L1 type knows that if she attacks, she will receive positive payoff D whenever $\theta \leq l$ and pay a fixed cost t with certainty. For a given l , the expected gain from attacking can be written as $D \cdot Pr(\theta \leq l|x, y) - t$. Given the conditional distribution of θ from (1), the probability of a successful attack is

¹⁰L0’s behavior may also depend upon other variables like D or t . The results presented below are robust to their inclusion.

$$Pr(\theta \leq l|x, y) = \Phi \left(\sqrt{\alpha + \beta} (l - (\mu x + (1 - \mu)y)) \right).$$

Averaging over l according to density q gives the expected payoff gain for an L1 type who receives signals x and y . This can be written as

$$\pi^{L1}(x) = D \int_0^1 \Phi \left(\sqrt{\alpha + \beta} (l - (\mu x + (1 - \mu)y)) \right) q(l|x, y) dl - t.$$

Note that π^{L1} is continuous and strictly decreasing in x by A1 and A2. There is a unique point κ_1 such that $\pi^{L1}(\kappa_1) = 0$. The cutoff κ_1 is determined implicitly by

$$\int_0^1 \Phi \left(\sqrt{\alpha + \beta} (l - (\mu x + (1 - \mu)y)) \right) q(l|x, y) dl = \frac{t}{D}. \quad (2)$$

L1 types play according to a threshold strategy with cutoff κ_1 . They attack if and only if their private signal x is below the threshold cutoff κ_1 .^{11,12}

L2 types best respond to the belief that all others are playing a threshold strategy with cutoff κ_1 . L2 types believe that players attack only if their private signal is below the threshold cutoff κ_1 . From the perspective of L2 types, there is a successful attack if and only if

$$\Phi \left(\sqrt{\beta} (\kappa_1 - \theta) \right) \geq \theta.$$

Let $\bar{\theta}^2$ be determined uniquely by the solution to $\Phi(\sqrt{\beta}(\kappa_1 - \bar{\theta}^2)) = \bar{\theta}^2$. As a result, an L2 type expects an attack to be successful whenever $\theta \leq \bar{\theta}^2$.

The expected payoff gain for an L2 type who observes signals x and y is given by $D \cdot Pr(\theta \leq \bar{\theta}^2|x, y) - t$. Using the conditional distribution of θ from (1), the expected payoff gain for an L2 type is given by

$$\pi^{L2}(x, \bar{\theta}^2) = D \cdot \Phi \left(\sqrt{\alpha + \beta} (\bar{\theta}^2 - (\mu x + (1 - \mu)y)) \right) - t.$$

Since π^{L2} is continuous and strictly decreasing in x , there is a unique point κ_2 such that $\pi^{L2}(\kappa_2, \bar{\theta}^2) = 0$. Therefore, L2 types play according to a threshold strategy with

¹¹Assumption A2 guarantees continuity in expected payoffs. However, without this assumption there still exists a unique point κ such that $\pi^{L1} > 0$ if $x < \kappa$ and $\pi^{L1} < 0$ if $x > \kappa$. As a result, L1 types still play threshold strategies, although they may no longer be indifferent between attacking and not attacking at the threshold signal.

¹²Also, note that this is for a given public signal y , thus $\kappa_1(y)$, but we suppress the notation throughout the paper.

cutoff κ_2 . They attack if and only if $x \leq \kappa_2$.

The behavior of higher types is similar to that of L2 types since they also believe that others are playing threshold strategies. Lk types will play a threshold strategy with cutoff κ_k , where κ_k is determined implicitly by

$$\Phi\left(\sqrt{\alpha + \beta}(\bar{\theta}^k - (\mu\kappa_k + (1 - \mu)y))\right) = \frac{t}{D} \quad (3)$$

with $\bar{\theta}^{Lk}$ determined by the solution to

$$\Phi\left(\sqrt{\beta}(\kappa_{k-1} - \bar{\theta}^k)\right) = \bar{\theta}^k. \quad (4)$$

Proposition 1 summarizes these results.

Proposition 1. *Let A1 and A2 hold. In the speculative attack game, an Lk type attacks if and only if $x \leq \kappa_k$ for all $k \geq 1$, where κ_1 is determined by equation (2) and κ_k ($k > 1$) is determined by equations (3) and (4).*

3.1 Properties of the Level-k Solution

This section establishes some of the properties of the level-k solution. Proposition 2 establishes a set of comparative static results for the level-k model. Proposition 3 establishes the relationship between equilibrium and level-k behavior: if there is a unique equilibrium, then the level-k cutoffs converge monotonically towards the equilibrium cutoff as k converges to infinity. And, Corollary 1 establishes the level-k results for two special information treatments: when there is only private information and when there is only public information.

Define the likelihood of a successful attack to be proportional to the size of the interval $[-\infty, \bar{\theta}]$, where $\bar{\theta}$ is the aggregate threshold for a successful attack under level-k thinking (i.e. there will be a currency crisis if $\theta \leq \bar{\theta}$ and no crisis otherwise).¹³ The intuitive comparative static results hold under level-k thinking. If the payoff to attacking (D) decreases, the likelihood of a successful attack decreases. If the cost of attacking (t) increases, the likelihood of a successful attack decreases. If public information (y) increases, then the likelihood of a successful attack crisis decreases. Proposition 2 formalizes the comparative static results.

¹³The proof of Proposition 2 establishes that there exists a unique $\bar{\theta}$ such that given the distribution of cognitive types there is a currency crisis if and only if $\theta \leq \bar{\theta}$.

Proposition 2. *Let A1 and A2 hold. The likelihood of a successful attack falls whenever D decreases, t increases or y increases.*

The comparative static predictions of the level- k model are consistent with the existing experimental results in coordinated attack games. Heinemann et al. (2004), Duffy and Ochs (2012), and Cornand (2006) all test how thresholds respond to changes in the safe payoff t . They find that the likelihood of a successful attack falls whenever t increases.

Proposition 3 establishes that the threshold cutoffs for each level- k type converge monotonically to the equilibrium threshold as k tends to infinity. Given the specification for L0, the cutoffs either monotonically increase or decrease towards \bar{x}^E depending on whether $\kappa_1 < \bar{x}^E$ or $\kappa_1 > \bar{x}^E$, respectively.

Proposition 3. *Let A1 and A2 hold and $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$. Let \bar{x}^E be the equilibrium cutoff and $\{\kappa_k\}$ be the set of level- k cutoffs. Then the threshold cutoffs $\{\kappa_k\}$ converge monotonically towards the equilibrium cutoff \bar{x}^E as $k \rightarrow \infty$. In addition,*

- (i) *if $\kappa_1 < \bar{x}^E$, then $\{\kappa_k\}$ is a strictly increasing sequence*
- (ii) *if $\kappa_1 > \bar{x}^E$, then $\{\kappa_k\}$ is a strictly decreasing sequence*

If the L1 cutoff is above the equilibrium cutoff under private information, then the L1 cutoff will be higher than the cutoffs of all other L k types (i.e. $\kappa_1 > \kappa_2 > \kappa_3 > \dots$). For this to hold we require the behavior of L0 types to be biased towards the payoff dominant equilibrium. This is a natural specification for L0. Many experiments find that players tend to play payoff-dominant actions in coordination games, at least in initial periods (e.g. Costa-Gomes et al., 2009).

For example, let $P = \{\frac{1}{2}, 1; \frac{1}{2}, 0\}$ be the distribution where half of the players attack with probability one. P represents the beliefs of the equilibrium threshold player when there is only private information. Because information is symmetric, each player believes that half the players should receive a signal above her own. Therefore, the threshold player believes that half the players will attack. If Q first-order stochastically dominates P , L0 behavior is biased towards the payoff-dominant equilibrium and we are guaranteed that all level- k thresholds will be above the equilibrium threshold.

The next corollary formalizes the results when there is only public or only private information.

Corollary 1. *Let A1 and A2 hold. Suppose players receive a private signal x_i or a public signal y with $x_i, y \sim N(\theta, \frac{1}{\beta})$. Let θ be distributed uniformly on the real line. There exists a unique cutoff, κ^Y , and unique cutoffs, κ_k^X , for each $k \geq 1$, such that*

- (i) Lk types attack if and only if $x_i \leq \kappa_k^X$, for $k \geq 1$ when information is private*
- (ii) Lk types attack if and only if $y \leq \kappa^Y$, for $k \geq 1$ when information is public*

When there is only public information, all types coordinate on the same threshold cutoff. However, when information is private, different types use different thresholds. This results from the interaction of bounded depths of reasoning with differences in the nature of public and private information. Public information coordinates higher- and lower-order beliefs about fundamentals. All types share the same beliefs about fundamentals, regardless of their depth of reasoning. This coordinates the behavior of players with different depths of reasoning. When information is private, higher-order beliefs are more dispersed than lower-order beliefs. This decreases coordination by causing differences in behavior for players with different depths of reasoning. The level-k model predicts greater coordination under public information than under private information.

3.2 Robustness

Threshold strategies, comparative statics, and the coordinating role of public information are robust predictions of limited depth of reasoning models. These results hold under only weak assumptions on the behavior of L0. And, they hold for any limited depth of reasoning model.

More specifically, the regularities only rely on assumption A1. A1 is a very weak requirement and holds trivially under the specification that L0 behavior does not depend on information. The level-k model also places specific assumptions on the beliefs of each cognitive type about the cognitive types of others. Specifically, an Lk type believes there are only Lk-1 types. Cognitive hierarchy, another popular limited depth of reasoning model, assumes that a Lk type believes there are all lower types with weights determined by a conditional Poisson distribution. In general, we could think that an Lk type may hold any distribution of beliefs over the lower types. The exact belief structure does not matter. Moreover, the exact type distribution does not matter. The results hold for any distribution of types in the population. In other words, they hold for any limited depth of reasoning model. This is shown in Appendix

B.

The results even hold for L_∞ types. L_∞ types have infinite depths of reasoning but believe others have bounded depths of reasoning. Further, the results hold for some proportion of sophisticated types. These are types that have infinite depths of reasoning and take into account the fact that others may also have infinite depths of reasoning. We show in Appendix B that as long as the proportion of sophisticated types is not too large, limited depth of reasoning with sophisticated types is consistent with the existing experimental evidence.

One of the main criticisms of the limited depth of reasoning literature is that predictions are not robust to model specification. This is not true for the aggregate predictions of the coordinated attack game. They are robust predictions of limited depth of reasoning models.

4 Empirical Analysis

This section analyzes the experimental data from Heinemann et al. (2004) (HNO hereafter).¹⁴ The distribution of level-k and equilibrium types are estimated using a finite mixture model. Allowing for level-k types significantly improves our ability to explain the experimental data.

For this section, the results rely on the specification of a particular limited depth of reasoning model. We use the level-k model. This is because the quantitative predictions of the model (the threshold cutoffs used by different types) are needed in order to fit the model to data. This does not mean that other limited depth of reasoning models would not fit the data equally well or even better than the level-k model. But the goal of this exercise is to show that there exists a limited depth of reasoning model that explains the data better than equilibrium. Notice that this exercise is different to that in Section 3, which concerns itself with qualitative predictions. In that case we showed limited depth of reasoning models make robust qualitative predictions that are consistent with existing experimental evidence.

¹⁴In this section, we restrict attention to the experimental data from Heinemann et al. (2004), because it includes both a private and public information treatment.

4.1 The experiment

Subjects play a finite player game similar to the above game. In each session, there are 15 subjects who simultaneously decide whether to attack or not attack. An attack is associated with an opportunity cost t (which is modeled as the safe payoff to not attacking). An attacking subject earns the amount θ if the attack is successful. An attack is successful if and only if a sufficient number of players attack. The threshold to a successful attack is determined by $a(\theta)$ which is a non-increasing function of θ . Low θ represents good fundamentals and a high threshold to a successful attack. High θ represents poor fundamentals and a low threshold to a successful attack.

This game differs from the game analyzed in previous sections in a number of ways: it is a discrete player game, the payoff from a successful attack depends upon the fundamentals, and a lower θ corresponds to better fundamentals. However, none of these changes alter the previous analysis in a substantial way.¹⁵ The only change to note is that the attack/non-attack regions are flipped. Players attack if and only if their signals are *above* the cutoff threshold.

There are two information treatments in the experiment. A private information treatment (PI) and a common information treatment (CI). In PI, players do not know the fundamental, but know that θ is distributed uniformly on $[10, 90]$ and receive private signals x_i randomly drawn with independent and uniform conditional distributions on $[\theta - 10, \theta + 10]$. In CI, players learn the fundamental θ when they receive a precise public signal $y = \theta$. There are also four different payoff parametrizations. The payoff to the safe (not attack) option varies between $t = 20$ and $t = 50$ and the threshold parameter ψ , $a(\theta) = \frac{15(80-\theta)}{\psi}$, varies between $\psi = 60$ and $\psi = 100$. The eight treatments are summarized in Table 1.

We analyze the data from a total of 29 sessions and 435 subjects. Subjects play 16 rounds in the experiment, but we analyze only the first round because we are interested in initial play. In each round, each subject receives 10 signals and makes a choice for each signal (attack or not-attack), for a total of 4350 choices.

¹⁵See Appendix C for an analysis of the level-k results. See HNO for an analysis of the equilibrium results.

Table 1: Treatments

Treatment	Information	Safe Payoff (t)	ψ	Sessions
PI_20_100	Private	20	100	1
CI_20_100	Public	20	100	1
PI_50_100	Private	50	100	1
CI_50_100	Public	50	100	1
PI_20_60	Private	20	60	6
CI_20_60	Public	20	60	5
PI_50_60	Private	50	60	7
CI_50_60	Public	50	60	7

4.2 Econometric Analysis

The analysis uses two main econometric methods. First, we follow Heinemann et al. (2004) and Duffy and Ochs (2012) and estimate aggregate mean thresholds by estimating a logit response model in which the binary attack decision depends on a constant and the signal ($Z = x$ or y). That is, we use maximum likelihood estimation to find the coefficient estimates \hat{a} and \hat{b} , that are a best fit to the logit response function:

$$Pr(\text{attack}|Z) = [1 + \exp(-a - bZ)]^{-1}$$

The attack threshold can be viewed as the critical value, Z^* , for which a representative player is indifferent between attacking and not attacking, which obtains when $Pr(\text{attack}|Z^*) = 0.5$. Using this, we can obtain the estimated mean attack threshold $\hat{Z}^* = -\frac{\hat{a}}{\hat{b}}$. The standard deviation is given by $\frac{\pi}{\hat{b}\sqrt{3}}$. We take the standard deviation to be a measure of the coordination of subjects around the estimated attack threshold.

Second, we follow Costa-Gomes et al. (2001) and estimate a finite mixture model that allows for different behavioral types. This approach assumes that each subject's type is drawn from a fixed common prior distribution over all types. The types we allow include: an equilibrium type (E) who believes all other types are equilibrium types, an L1 type who believes all other types are L0 types whose behavior is fixed (and specified below), and an L2 type who believes all other types are L1 types.

We assume that each player follows the predictions of a particular type with error. Because subjects often make a type's exact choices, we use a simple spike-logit error structure. Index types $k = 1, \dots, 3$ and choices by $q = 1, \dots, 10$. In each choice, a subject has a given probability $1 - \epsilon$ of making her type's exact choice, and with error-

rate, ϵ , makes choices that follow a logistic distribution with error density, $d_q^k(a_q^i, \lambda)$.¹⁶ The parameter λ represents the logistic precision parameter.

For subject i , let Q^{ik} represent the set of choices where subject i 's action is consistent with type k 's predicted action. The aggregate log-likelihood function is given by

$$\ln L(p, \epsilon, \lambda | a) = \sum_{i=1}^S \ln \left[\sum_{k=1}^3 p_k \left(\prod_{q \in Q^i} (1 - \epsilon + \epsilon d_q^k(a_q^i, \lambda)) \right) \left(\prod_{q \in Q^{iC}} \epsilon d_q^k(a_q^i, \lambda) \right) \right] \quad (5)$$

where S is the number of subjects.

With three types, this model has 4 independent parameters: 2 independent type probabilities p_k , a precision parameter λ , and an error-rate ϵ .¹⁷

In order to fully specify this model, we must specify the behavior of L0 types. In most level- k applications, L0 behavior is chosen to be either uniformly random or some focal behavior. However, we know from previous coordination experiments that players tend to focus on payoff-dominant outcomes (at least during initial play) (Costa-Gomes et al. 2009) and Crawford et al. (2008) find that level- k models with L0 players biased towards payoff-salient outcomes explain behavior in coordination games fairly well. For these reasons, we assume that the behavior of L0 types is given by a discrete probability function $q(l)$ over $\{0, 1, \dots, 14\}$ that first-order stochastically dominates the uniformly random distribution.¹⁸ This biases the behavior of L0 types towards the payoff-dominant outcome (attack) relative to the uniformly random specification. Notice that this specification of L0 behavior satisfies assumption A1.

¹⁶The error density, $d_q^k(a_q^i, \lambda)$, is defined $d_q^k(a_q^i, \lambda) = \frac{\exp[\lambda S_q^k(a_q^i)]}{\exp[\lambda S_q^k(attack)] + \exp[\lambda S_q^k(not-attack)]}$. The term, $S_q^k(a_q^i)$, is type k 's expected payoff from playing action a_q^i in choice q , given type k 's beliefs about the actions of others.

¹⁷We limit attention to an aggregate analysis because of the way the data is structured. Subjects make choices for 10 random signal draws. Depending on the random draws, we may not be able to identify a type for every subject. As well, the design is a between subjects design, subjects do not make choices under both public and private information. It is only in the private information treatments that L1 and L2 types are identified separately (although public information treatments are still helpful for separating level- k and equilibrium types). Thus, we use the aggregate analysis as a substitute for the individual analysis. Costa-Gomes et al. (2001) run both individual and aggregate analyses and find virtually the same results (individual results recorded in earlier version of the paper (Costa-Gomes et al. 1998)).

¹⁸Specifically, the distribution used is $q(0) = q(1) = q(2) = q(3) = q(4) = 1/60$, $q(5) = q(6) = q(7) = q(8) = q(9) = q(13) = q(14) = 4/60$, and $q(10) = q(11) = q(12) = 8/60$.

Each type’s behavior is determined by a threshold cutoff. A type attacks if and only if her signal is above the threshold cutoff. Given the behavior of the L0 types we can determine the threshold cutoffs for L1 and L2 types. The threshold cutoffs for equilibrium types are given by the unique equilibrium thresholds in the private information treatments and by the global games thresholds in the public information treatments.¹⁹ Thresholds vary across both information treatments and safe-payoff treatments for all four types. The type’s threshold predictions are given in Table 2. These thresholds fully specify each type’s predicted actions. A type k attacks if and only if her signal in choice q is greater than or equal to her cutoff threshold.

Table 2: Threshold Predictions by Type

		CI		PI		
t	ψ	E	L1, L2	E	L1	L2
20	100	33.3	27.3	32.4	27.7	29.5
50	100	60	53.3	61	53.4	58
20	60	44	37.5	41.8	38.3	39.6
50	60	64	57.7	66.0	57.9	61.2

4.3 Aggregate Regularities

The mean threshold cutoffs are estimated for each treatment using the binary response logit model specified in the previous section. Table 3 lists the estimated mean cutoff thresholds and the associated standard deviations for each treatment.²⁰

¹⁹HNO find that the global games equilibrium explains this data better than any other known equilibrium selection method, i.e. payoff-dominance, risk dominance or max-min.

²⁰The patterns we observe in the first round data are similar to the patterns observed by HNO across the entire 16 rounds.

Table 3: Estimated Mean Threshold Cutoffs

(n=4350)	Treatment							
	PI_20_100	CI_20_100	PI_50_100	CI_50_100	PI_20_60	CI_20_60	PI_50_60	CI_50_60
Estimated mean threshold	30.36	37.24	58.14	53.50	48.54	36.90	60.62	54.63
Estimated standard deviation	29.63	22.79	17.06	15.12	29.14	29.02	22.77	17.79

We can see from Tables 2 and 3 that mean thresholds are closer to the level-k predictions than the equilibrium predictions. As well, mean thresholds are lower in the PI20 and CI20 treatments relative to the PI50 and CI50 treatments. Subjects respond to changes in the safe payoff t as predicted in both the level-k and equilibrium models. Mean thresholds are also lower in the CI treatments relative to the PI treatments in all cases except when $t = 20$ and $\psi = 100$. The level-k model is consistent with these patterns. However, equilibrium predicts that the mean threshold should increase in the CI_50_60 treatment relative to PI_50_60, which is not true in the data.

Lastly, notice that the standard deviation of estimated thresholds is lower in the CI treatments relative to the corresponding PI treatments. This is inconsistent with an equilibrium analysis, which predicts that coordination should be weakly better under private information than public information. The level-k model is consistent with this pattern, as the level-k model predicts that coordination should be strictly better under public information than under private information.

4.4 Model Estimation

The type classification is based on 435 subjects who each make 10 choices. There are 279 choices that separate the predictions of E and L1 types, 222 choices that separate the predictions of E and L2 types and 57 choices that separate the predictions of L1 and L2 types.

Table 4 gives the maximum likelihood estimates of equation (5). The first column gives the estimates for the combined model that includes all three types: E, L1, and L2. The second column gives the estimates for the level-k model and the third column gives the estimates for the equilibrium model. In the combined model, level-k types make up 83 percent of the estimated type distribution. L1 types are most frequent, making up 70 percent of the type distribution and L2 types make up 13 percent. This

Table 4: Aggregate Type Classification

MODEL:	All Types	Level-k	Equilibrium
Log-Likelihood	-1955.80	-1960.16	-2138.97
L1	.6993 (.0487)	.6695 (.0542)	
L2	.1312 (.0556)	.3305 (.0542)	
E	.1695 (.0700)		1
λ	.0410 (.0032)	.0430 (.0028)	.0255 (.0020)
ϵ	.7366 (.0365)	.7906 (.0367)	.7084 (.0354)
n	4350	4350	4350
AIC	3953.5	3953.8	4294.7
BIC	3921.6	3928.4	4281.9

*Notes: bracketed numbers are bootstrapped standard errors clustered at the subject level.

is compared to 17 percent for equilibrium types.

The level-k model fits the data better than the equilibrium model looking at both AIC and BIC. But, the combined model does not really offer an improvement over the level-k model. The AIC for the model with all four types is 3953.5, which is only marginally smaller than the AIC for the level-k model, 3953.8.

Both the precision parameter and the error-rate parameter combine to determine the rate of deviation from each type's predicted play. Types follow their type's predicted action about 83 percent of the time on average.

5 Conclusion

The existing experimental evidence on coordinated attack games presents a challenge to equilibrium theory, which mispredicts the coordinating roles of public and private

information. Limited depth of reasoning models provide a consistent explanation of experimental behavior across different information conditions. Having a model that explains experimental behavior under private and public information provides a behaviorally-motivated answer to the debate over the potentially destabilizing effects of public information.

In addition, this paper highlights robust predictions of the level-k model. The main results from coordinated attack experiments are robust predictions of limited depth of reasoning models. This demonstrates that limited depth of reasoning models have predictive power even if L0 and the type distribution are not stable across games. In particular, comparative static predictions hold for a wide class of these models. Determining accurate comparative statics does not require determining accurate point predictions. This is especially desirable if what we care about is policy prescriptions and not necessarily point estimates.

References

- Bannier, C. (2002). Private and Public Information in a Self-fulfilling Currency Crisis. *Journal of Economics*, 76(1):65–85.
- Brocas, I., Carrillo, J. D., Wang, S. W., and Camerer, C. F. (2014). Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games. *Review of Economic Studies*, 81(3):944–970.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119(3):861–898.
- Carlsson, H. and van Damme, E. (1993a). *Equilibrium Selection in Stag Hunt Games*, pages 237–253. MIT Press, Cambridge.
- Carlsson, H. and van Damme, E. (1993b). Global Games and Equilibrium Selection. *Econometrica*, 61(5):989–1018.
- Cornand, C. (2006). Speculative Attacks and Information Structure: An Experimental Study. *Review of International Economics*, 14(5):797–817.
- Cornand, C. and Heinemann, F. (2014). Measuring agents’ reaction to private and

- public information in games with strategic complementarities. *Experimental Economics*, 17:61–77.
- Cornand, C. and Heinemann, F. (2015). Limited higher order beliefs and the welfare effects of public information. *Journal of Economic Studies*, 42(6):1–26.
- Costa-Gomes, M. and Crawford, V. P. (2006). Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. *American Economic Review*, 96(5):1737–1768.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (1998). Cognition and Behavior in Normal-Form Games: An Experimental Study. *UCSD Discussion Paper 98-22*.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69(5):1193–1235.
- Costa-Gomes, M. A., Crawford, V. P., and Iriberri, N. (2009). Comparing Models of Strategic Thinking in Van Huyck, Battalio and Beil’s Coordination Games. *Journal of the European Economic Association Papers and Proceedings*, 7(2):365–376.
- Costa-Gomes, M. A., Crawford, V. P., and Iriberri, N. (2013). Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature*, 51:5–62.
- Crawford, V. P., Gneezy, U., and Rottentreich, Y. (2008). The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures. *American Economic Review*, 98(4):1443–1458.
- Crawford, V. P. and Iriberri, N. (2007). Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions. *Econometrica*, 75(6):1721–1770.
- Duffy, J. and Ochs, J. (2012). Equilibrium Selection in Static and Dynamic Entry Games. *Games and Economic Behavior*, 76:97–116.
- Hartigan, J. (1983). *Bayes Theory*. Springer-Verlag, New York.
- Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information. *Econometrica*, 72(5):1583–1599.

- Heinemann, F., Nagel, R., and Ockenfels, P. (2009). Measuring Strategic Uncertainty in Coordination Games. *Review of Economic Studies*, 76(1):181–221.
- Hellwig, C. (2002). Imperfect Common Knowledge of Preferences in Global Coordination Games. *mimeo*.
- Kneeland, T. (2012). Global Games under Limited Depth of Reasoning. *working paper*.
- Morris, S. and Shin, H. S. (1998). Unique Equilibrium in a Model of Self-fulfilling Attacks. *American Economic Review*, 88(3):587–597.
- Morris, S. and Shin, H. S. (2003). *Global Games: Theory and Applications*, pages 56–114. Cambridge University Press, Cambridge.
- Morris, S. and Shin, H. S. (2004). Coordination Risk and the Price of Debt. *European Economic Review*, 48(1):133–153.
- Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review*, 85(5):1313–1326.
- Rubinstein, A. (1989). The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge". *American Economic Review*, 79(3):385–391.
- Shapiro, D., Shi, X., and Zillante, A. (2014). Level-k Reasoning in a Generalized Beauty-Contest. *Games and Economic Behavior*, 86:308–329.
- Stahl, D. O. and Wilson, P. W. (1994). Experimental Evidence on Player's Models of Other Players. *Journal of Economic Behavior and Organization*, 25(3):309–327.
- Stahl, D. O. and Wilson, P. W. (1995). On Player's Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254.
- Strzalecki, T. (2014). Depth of Reasoning and Higher Order Beliefs. *Journal of Economic Behavior and Organization*, 108:108–122.

A Omitted Proofs

Proof of Proposition 1

The expected payoff gain for an L1 type who receives signal x can be written as

$$\pi^{L1}(x) = D \int_{-\infty}^{\infty} \phi\left(\sqrt{\alpha + \beta}(\theta - \mu x - (1 - \mu)y)\right) [1 - Q(\theta|x, y)] d\theta - t$$

This expression is continuous and strictly decreasing in x by A1 and A2. To see this notice that there are two effects to consider. First, $1 - Q(\theta|x, y)$ is weakly decreasing in x for a given θ . Second, the mean of the distribution of θ shifts as x increases which puts more weight on higher values of θ , which means more weight is put on smaller values of $1 - Q(\theta|x, y)$. The two effects work in the same direction. Also, $\pi^{L1}(-\infty) > 0$ and $\pi^{L1}(\infty) < 0$ by properties of the normal distribution. As a result, L1 types play a threshold strategy with cutoff κ_1 determined implicitly by equation (2).

Claim: Lk types play threshold strategies with cutoff $\kappa_k \forall k \geq 1$.

This claim is proved by induction on k .

It is true for L1 types by the above argument.

Suppose it is true for Lk-1 types.

Lk types expect a proportion $\Phi(\sqrt{\beta}(\kappa_{k-1} - \theta))$ of players to attack. Since Φ is continuous and strictly decreasing in θ , we can find a unique $\bar{\theta}^k$ such that an Lk type thinks there is a successful attack if and only if $\theta \leq \bar{\theta}^k$ where $\bar{\theta}^k$ is determined implicitly by equation (4).

The expected payoff gain for an Lk type who receives signal x is then given by

$$\pi^{Lk}(x, \bar{\theta}^k) = D \left(1 - \Phi(\sqrt{\alpha + \beta}(\mu x + (1 - \mu)y - \bar{\theta}^k)) \right) - t.$$

π^{Lk} is continuous, strictly decreasing in x , $\pi^{Lk}(-\infty) > 0$ and $\pi^{Lk}(\infty) < 0$. As a result, Lk types play according to a threshold strategy with cutoff κ_k determined by implicitly by equation (3).

The result follows by induction.

Proof of Proposition 2

Consider the case for level- k types. Let the distribution of types be specified by $\lambda_1, \lambda_2, \dots$ where λ_k is the proportion of Lk types.

Let $LK(\theta)$ be the proportion of all Lk types that attack,

$$LK(\theta) = \sum_{k=1}^{\infty} \lambda_k \Phi \left(\sqrt{\beta}(\kappa_k - \theta) \right)$$

$LK(\theta)$ is continuous and strictly decreasing in θ . Thus, there exists a unique $\bar{\theta}$ such that $LK(\bar{\theta}) = \bar{\theta}$ and there is a currency crisis iff $\theta \leq \bar{\theta}$.

We want to know how $\bar{\theta}$ changes when some parameter z varies. If $\frac{\partial \kappa_k}{\partial z} > 0 (< 0) \forall k$, then $LK(\theta)$ is shifted to the right (left) and hence $\frac{\partial \bar{\theta}}{\partial z} > 0 (< 0)$.

Claim: $\frac{\partial \kappa_k}{\partial D} > 0$, $\frac{\partial \kappa_k}{\partial t} < 0$ and $\frac{\partial \kappa_k}{\partial y} < 0 \forall k \geq 1$

Show by induction on k .

Consider $k = 1$. Let

$$F = D \int_0^1 \left(1 - \Phi \left(\sqrt{\alpha + \beta}(\mu \kappa_1 + (1 - \mu)y - l) \right) \right) q(l|\kappa_1, y) dl - t$$

The cutoff is implicitly determined by $F = 0$.

By the proof of Proposition 1 $\frac{\partial F}{\partial \kappa_1} < 0$.

$$\frac{\partial F}{\partial D} = \int_0^1 \left(1 - \Phi \left(\sqrt{\alpha + \beta}(\mu \kappa_1 + (1 - \mu)y - l) \right) \right) q(l|\kappa_1, y) dl > 0 ,$$

$$\frac{\partial F}{\partial t} = -1 < 0 ,$$

$$\frac{\partial F}{\partial y} < 0. \text{ This is analogous to } \frac{\partial F}{\partial x}.$$

Therefore, by the implicit function theorem, $\frac{\partial \kappa_1}{\partial D} > 0$, $\frac{\partial \kappa_1}{\partial t} < 0$ and $\frac{\partial \kappa_1}{\partial y} < 0$.

Now suppose the claim is true for $k - 1$.

Equations (3) and (4) jointly determine κ_k .

From the equation (4) we have:

$$\frac{d\bar{\theta}^{Lk}}{dz} = \frac{\sqrt{\beta}\phi(\cdot)}{1 + \sqrt{\beta}\phi(\cdot)} \frac{d\kappa_{k-1}}{dz}$$

Totally differentiate equation (3) taking into account the effect on the cutoff $\bar{\theta}^{Lk}$:

$$\left(\sqrt{\alpha + \beta}\mu\phi(\cdot) \right) d\kappa_k - \left(\frac{t}{D^2} + \sqrt{\alpha + \beta}\phi(\cdot) \frac{d\bar{\theta}^{Lk}}{dD} \right) dD = 0 \Rightarrow \frac{\partial \kappa_k}{\partial D} > 0$$

$$\left(\sqrt{\alpha + \beta}\mu\phi(\cdot) \right) d\kappa_k + \left(\frac{1}{D} - \sqrt{\alpha + \beta}\phi(\cdot) \frac{d\bar{\theta}^{Lk}}{dt} \right) dt = 0 \Rightarrow \frac{\partial \kappa_k}{\partial t} < 0$$

$$\left(\sqrt{\alpha + \beta}\mu\phi(\cdot) \right) d\kappa_k + \sqrt{\alpha + \beta}\phi(\cdot) \left((1 - \mu) - \frac{d\bar{\theta}^{Lk}}{dy} \right) dy = 0 \Rightarrow \frac{\partial \kappa_k}{\partial y} < 0$$

The claim follows by induction.

Thus, by the argument above: $\frac{\partial \bar{\theta}}{\partial D} > 0$, $\frac{\partial \bar{\theta}}{\partial t} < 0$ and $\frac{\partial \bar{\theta}}{\partial y} < 0$.

Proof of Proposition 3

Define the function $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$f(x, \kappa) = D \left(1 - \Phi(\sqrt{\alpha + \beta}(\mu x + (1 - \mu)y - \bar{\theta}(\kappa))) \right) - t$$

where $\bar{\theta}(\kappa)$ is implicitly determined by $\bar{\theta}(\kappa) = \Phi(\sqrt{\beta}(\kappa - \bar{\theta}(\kappa)))$.

Define the function $b : \mathbb{R} \rightarrow \mathbb{R}$ by letting $b(\kappa)$ be the value that solves $f(b(\kappa), \kappa) = 0$.

By argument in proof of Proposition 1, b is a well-defined function. By the implicit function theorem $\frac{\partial \bar{\theta}}{\partial \kappa} > 0$ and $\frac{\partial b}{\partial \kappa} > 0$. So, b is a strictly increasing, continuous function. Notice that $f(\bar{x}, \bar{x}) = 0$ defines the equilibrium cutoff \bar{x} . This means that b has a unique fixed point. Define the function $b^k : \mathbb{R} \rightarrow \mathbb{R}$ recursively by $b^k(\kappa) = b(b^{k-1}(\kappa))$ for $k \geq 2$.

Now let κ_1 be such that $\kappa_1 < b(\kappa_1)$.

Consider the following real sequence $\{b^n(\kappa_1)\}_{n=1}^{\infty}$. This sequence is increasing, since b is an increasing function and $\kappa_1 < b(\kappa_1)$. The sequence is also bounded since b is a bounded function (bounded above by \bar{x})

Since $\{b^n(\kappa_1)\}$ is bounded, continuous, and increasing, it has some limit, call it c .

$$c = \lim_{n \rightarrow \infty} b^{n+1}(\kappa_1) = \lim_{n \rightarrow \infty} b(b^n(\kappa_1)) = b\left(\lim_{n \rightarrow \infty} b^n(\kappa_1)\right) = b(c)$$

where the third equality follows from continuity of b . Since, b has a unique fixed point $c = \bar{x}$.

We can find the analogous result for any κ_1 such that $b(\kappa_1) < \kappa_1$.

The proof follows by noticing that $\kappa_k = b^{k-1}(\kappa_1)$.

Proof of Corollary 1

Let z be the signal received (i.e.. either $z = x$ or $z = y$).

The expected payoff gain of an L1 type is given by

$$\pi_z^{L1}(z) = D \int_{-\infty}^{\infty} \phi(\sqrt{\beta}(\theta - z))[1 - Q(\theta|z)]d\theta - t.$$

By similar arguments as in Proposition 1, π^{L1} is strictly decreasing and continuous in z . Therefore, L1 types play threshold strategies with cutoff κ_1^z determined by $\pi_z^{L1}(\kappa_1^z) = 0$. Players attack if signal $z \leq \kappa_1^z$ and do not attack otherwise.

Claim1: Lk types play threshold strategies with cutoff $\kappa_k^X \forall k \geq 1$ in the private information game

This claim is proved by induction on k.

It is true for L1 types by the above argument.

Suppose it is true for Lk-1 types.

Lk types expect a proportion $\Phi(\sqrt{\beta}(\kappa_{k-1}^X - \theta))$ of players to attack. Since Φ is continuous and strictly decreasing in θ , we can find a unique $\bar{\theta}^k$ such that an Lk type thinks there is a successful attack if and only if $\theta \leq \bar{\theta}^k$ where $\bar{\theta}^k$ is determined implicitly by $\Phi(\sqrt{\beta}(\kappa_{k-1}^X - \bar{\theta}^k)) = \bar{\theta}^k$.

The expected payoff gain for an Lk type who receives signal x is then given by

$$\pi^{Lk}(x, \bar{\theta}^k) = D \left(1 - \Phi(\sqrt{\beta}(z - \bar{\theta}^k)) \right) - t.$$

π^{Lk} is continuous, strictly decreasing in x , $\pi^{Lk}(-\infty) > 0$ and $\pi^{Lk}(\infty) < 0$. As a result, Lk types play according to a threshold strategy with cutoff κ_k^X uniquely determined by $\pi^{Lk}(\kappa_k^X, \bar{\theta}^k) = 0$

The result follows by induction.

Claim2: Lk types play threshold strategies with cutoff $\kappa^Y \forall k \geq 1$ in the public information game

This claim is proved by induction on k.

It is true for L1 types by the above argument.

Suppose it is true for Lk-1 types.

Lk types expect everyone to attack if $y \leq \kappa^Y$.

If $y > \kappa^Y$, payoff to attacking is $D \cdot Pr(\theta \leq 0) - t$ (< 0 since L1 payoff is negative)

If $y \leq \kappa^Y$ then payoff to attacking is $D \cdot Pr(\theta \leq 1) - t$ (> 0 since L1 payoff is positive)

The result follows by induction.

B

Generalized Level-k

The extension from level-k thinking to a general limited depth of reasoning model follows almost immediately. L1 types best respond only to L0. As a result, the behavior of an L1 type is unchanged under any limited depth of reasoning model.

Suppose L1 types play a threshold strategy with cutoff κ_1 . Let λ_i^k be the proportion of Li types that an Lk type believes she is playing against.

Claim: Lk types play threshold strategies with cutoff $\kappa_k \forall k \geq 1$.

This claim is proved by induction on k.

It is true for L1 types trivially, analysis doesn't change.

Suppose it is true for Lk-1 types.

First consider the case where $\lambda_0^k \neq 0$. Lk types believe the proportion of players attacking is given by $\lambda_0^k l_0 + \lambda_1^k \Phi(\sqrt{\beta}(\kappa_1 - \theta)) + \dots + \lambda_{k-1}^k \Phi(\sqrt{\beta}(\kappa_{k-1} - \theta))$. Thus, Lk types believe they get a payoff D if $\lambda_0^k l_0 + \lambda_1^k \Phi(\sqrt{\beta}(\kappa_1 - \theta)) + \dots + \lambda_{k-1}^k \Phi(\sqrt{\beta}(\kappa_{k-1} - \theta)) \geq \theta$. Using the distribution Q , we can then write the expected payoff of an Lk type as

$$\pi^{Lk}(x, \kappa_1, \dots, \kappa_{k-1}) = \int_{-\infty}^{\infty} \phi(\sqrt{\alpha + \beta}(\mu x + (1-\mu)y - \theta)) \left(1 - Q \left[\frac{1}{\lambda_0^k} \left(\theta - \sum_{i=1}^{k-1} \lambda_i^k \Phi(\sqrt{\beta}(\kappa_i - \theta)) \right) |x, y \right] \right) d\theta$$

π^{Lk} is continuous, strictly decreasing in x by A1 and A2 and $\pi^{Lk}(-\infty) > 0$ and $\pi^{Lk}(\infty) < 0$. To see that π^{Lk} is strictly decreasing notice that $1 - Q$ is weakly decreasing in x for a given θ . As x increases, the mean of the normal distribution shifts, placing more weight on higher values of θ . So more weight is placed on lower values of $1 - Q$. The two effects work in the same directions. As a result, Lk types play a threshold strategy with a cutoff κ_k where κ_k is uniquely determined by $\pi_{\sigma}^{Lk}(\kappa_k, \kappa_1, \dots, \kappa_{k-1}) = 0$.

Now consider the other case where $\lambda_0^k = 0$. Lk types believe the proportion of players attacking is given by $\lambda_1^k \Phi(\sqrt{\beta}(\kappa_1 - \theta)) + \dots + \lambda_{k-1}^k \Phi(\sqrt{\beta}(\kappa_{k-1} - \theta))$. Thus, Lk types believe an attack is successful if $\lambda_1^k \Phi(\sqrt{\beta}(\kappa_1 - \theta)) + \dots + \lambda_{k-1}^k \Phi(\sqrt{\beta}(\kappa_{k-1} - \theta)) \geq \theta$. The LHS is continuous and strictly decreasing in θ . Therefore \exists a unique $\bar{\theta}^{Lk}$ such that there is a successful attack iff $\theta \leq \bar{\theta}^{Lk}$. The expected payoff for an Lk type is then given by

$$\pi^{Lk}(x, \bar{\theta}^{Lk}) = D \left(1 - \Phi(\sqrt{\alpha + \beta}(\mu x + (1-\mu)y - \bar{\theta}^{Lk})) \right) - t.$$

π^{Lk} is continuous, strictly decreasing in x , $\pi^{Lk}(-\infty) > 0$ and $\pi^{Lk}(\infty) < 0$. As a result, Lk types play according to a threshold strategy with cutoff κ_k uniquely determined by $\pi^{Lk}(\kappa_k, \bar{\theta}^{Lk}) = 0$.

The result follows by induction.

Sophisticated Types

Proposition. *Let γ be the proportion of sophisticated players and let A1 and A2 hold. If $\gamma \frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$, there is a unique equilibrium in the speculative attack game where sophisticated types attack if and only if $x \leq \kappa^s$, where κ^s is uniquely determined.*

PROOF:

Consider the optimal strategy for a sophisticated type who believes all other sophisticated types are playing according to a threshold strategy with cutoff κ_s . Let the distribution of types be specified by $\lambda_1, \lambda_2, \dots$ where λ_k is the proportion of Lk types. A type Lk plays according to a threshold strategy with cutoff κ_k . The proportion of players attacking is given by

$$L(\theta) = (1 - \gamma) \sum_{k=1}^{\infty} \lambda_k \Phi\left(\sqrt{\beta}(\kappa_k - \theta)\right) + \gamma \Phi(\sqrt{\beta}(\kappa_s - \theta))$$

There is a unique value $\bar{\theta}^s$ such that $L(\bar{\theta}^s) = \bar{\theta}^s$. There is a currency crisis if and only if $\theta \leq \bar{\theta}^s$. Rearranging the expression, $L(\bar{\theta}^s) = \bar{\theta}^s$, and letting $S(\theta) = \sum_{k=1}^{\infty} \lambda_k \Phi(\sqrt{\beta}(\kappa_k - \theta))$, we can get the following relationship between κ_s and $\bar{\theta}^s$,

$$\kappa_s = \bar{\theta}^s + \frac{1}{\sqrt{\beta}} \Phi^{-1}\left(\frac{1}{\gamma} \bar{\theta}^s - \frac{1 - \gamma}{\gamma} S(\bar{\theta}^s)\right).$$

An equilibrium cutoff κ_s is the solution to the following equation

$$D \left(1 - \Phi(\sqrt{\alpha + \beta}(\mu \kappa_s + (1 - \mu)y - \bar{\theta}^s))\right) = t.$$

The aggregate size of attack $\bar{\theta}^s$ is then determined implicitly by equation

$$\frac{\alpha}{\sqrt{\beta}}(y - \bar{\theta}^s) + \Phi^{-1}\left(\frac{1}{\gamma} \bar{\theta}^s - \frac{1 - \gamma}{\gamma} S(\bar{\theta}^s)\right) = \frac{\sqrt{\alpha + \beta}}{\sqrt{\beta}} \Phi^{-1}\left(\frac{D - t}{D}\right).$$

The left hand side of this expression is continuous in $\bar{\theta}^s$, positive for low $\bar{\theta}^s$ and negative for high $\bar{\theta}^s$. As a result, there always exists a solution. In other words, there exists at least one threshold equilibrium. There will be a unique symmetric threshold equilibrium whenever there is a unique solution $\bar{\theta}^s$ to this equation.

$$\text{Let } G(\theta) = \frac{\alpha}{\sqrt{\beta}}(y - \theta) + \Phi^{-1}\left(\frac{1}{\gamma} \theta - \frac{1 - \gamma}{\gamma} S(\theta)\right).$$

$$\text{Then, } \frac{dG}{d\theta} = \frac{\frac{1}{\gamma} + \frac{1 - \gamma}{\gamma} (-S'(\theta))}{\phi(\Phi^{-1}(\frac{1}{\gamma} \theta - \frac{1 - \gamma}{\gamma} S(\theta)))} - \frac{\alpha}{\sqrt{\beta}}.$$

Since $\phi \leq 1/\sqrt{2\pi}$ and $-S'(\theta) > 0$ hold, G is strictly increasing whenever $\sqrt{2\pi} \geq \frac{\alpha}{\sqrt{\beta}}\gamma$.

Uniqueness follows from an iterated deletion of strictly dominated strategies argument. Let $\pi^s(x, \kappa)$ be the expected gain for a sophisticated player attacking, given that she observes signal x and all other sophisticated players are playing threshold strategies with a cutoff κ .

$$\pi^s(x, \kappa) = D \left(1 - \Phi(\sqrt{\alpha + \beta}(\mu x + (1 - \mu)y - \bar{\theta}^s(\kappa))) - t \right)$$

Define the function $b : \mathbb{R} \rightarrow \mathbb{R}$ to be the value $b(\kappa)$ that solves the equation $\pi^s(b(\kappa), \kappa) = 0$. b is well-defined, continuous and strictly increasing by arguments above. Therefore, $b(\kappa)$ has a unique fixed point, κ^s . Define the function $b^k : \mathbb{R} \rightarrow \mathbb{R}$ recursively by $b^k(\kappa) = b(b^{k-1}(\kappa))$ for $k \geq 2$.

Consider strategies that survive 1 round of deletion of strictly dominated strategies. First, consider strategies that are dominated by playing attack at some signals. The best payoff that a player could achieve if not attacking at those signals is if all sophisticated types played not-attack for all signals. That is, as if all sophisticated types played the threshold strategy with cutoff $-\infty$. The best response to sophisticated types playing with a cutoff $-\infty$ is to play a threshold strategy with a cutoff $b(-\infty)$. This means that any strategy for which not-attack is played for any $x < b(-\infty)$ is a dominated strategy. Following an analogous argument we can show that any strategy for which attack is played for any $x > b(\infty)$ is a dominated strategy. Thus, the set of strategies that survive 1 round of iterated deletion of strictly dominated strategies look like:

$$\text{play } \begin{cases} \text{attack} & \text{if } x < b(-\infty) \\ \text{not-attack} & \text{if } x > b(\infty) \end{cases}$$

We can now repeat the argument on this smaller set of strategies and show that the set of strategies that survive k -rounds of iterated deletion of strictly dominated strategies look like:

$$\text{play } \begin{cases} \text{attack} & \text{if } x < b^k(-\infty) \\ \text{not-attack} & \text{if } x > b^k(\infty) \end{cases}$$

By the argument in the proof of Proposition 3 for a function b that is strictly increasing, continuous, bounded and has one fixed point: $b^n(x_0) \rightarrow \kappa^s$ as $n \rightarrow \infty$ for any $x_0 \in \mathbb{R}$. This means that there is a unique strategy that survives iterated deletion of strictly dominated strategies - the threshold strategy with cutoff κ^s .

C

Proposition. *Suppose the behavior of L0 types is given by the cdf $Q(l)$ on $\{0, \dots, 14\}$.*

(i) *Suppose we are in information treatment PI. Then there exists a unique κ_k , for each $k \geq 1$, such that the behavior of type Lk is given by*

$$\text{play} \begin{cases} 1 & \text{if } x \geq \kappa_k \\ 0 & \text{o.w} \end{cases}$$

(ii) *Suppose we are in information treatment CI. Then there exists a unique κ such that the behavior of type Lk, for any $k \geq 1$, is given by*

$$\text{play} \begin{cases} 1 & \text{if } y \geq \kappa \\ 0 & \text{o.w} \end{cases}$$

PROOF

(i) By induction on k.

Consider k=1.

The expected payoff gain from playing action 1 for an L1 type who receives a private signal x , given the behavior of L0 is

$$\begin{aligned} \pi^{L1}(x) &= \frac{1}{2\epsilon} \int_{x-10}^{x+10} \int_{\frac{a(\theta)-1}{15}}^1 \theta q(l) dl d\theta - t \\ &= \frac{1}{2\epsilon} \int_{x-10}^{x+10} \theta \left[1 - Q\left(\frac{a(\theta)-1}{15}\right) \right] d\theta - t \end{aligned}$$

Notice that π^{L1} is strictly increasing and continuous in x . Also, for θ low enough $1 - Q = 0$ and for θ high enough, $1 - Q = 1$. Thus, there exists x_l, x_h such that $\pi^{L1}(x_l) < 0$ and $\pi^{L1}(x_h) > 0$. Therefore, L1 types play according to a threshold strategy with cutoff κ_1 determined by $\pi^{L1}(\kappa_1) = 0$.

Assume true for k-1.

Consider k.

The expected payoff gain of an Lk type given that she receives signal x and that she expects all other players to play threshold strategies with the cutoff κ_{k-1} is

$$\pi^{Lk}(x, \kappa_{k-1}) = \frac{1}{2\epsilon} \int_{x-10}^{x+10} \theta \left[1 - \text{Bin}\left(\hat{a}(\theta) - 2, n - 1, \frac{\theta + \epsilon - \kappa_{k-1}}{2\epsilon}\right) \right] d\theta - t$$

The function π^{Lk} is strictly increasing and continuous in x and there exists x_l, x_h such that $\pi^{Lk}(x_l) < 0$ and $\pi^{Lk}(x_h) > 0$. Therefore, there is a unique cutoff κ_k where

$\pi(\kappa_k, \kappa_{k-1}) = 0$. Best responses for Lk types are to play threshold strategies with cutoff κ_k .

The result follows by induction.

(ii) By induction on k.

Consider k=1.

The expected payoff gain from playing action 1 for an L1 type who receives a precise public signal y is

$$\begin{aligned}\pi^{L1}(y) &= \int_{\frac{a(y)-1}{15}}^1 yq(l)dl - t \\ &= y \cdot \left[1 - Q\left(\frac{a(y)-1}{15}\right) \right] - t\end{aligned}$$

Notice that π^{L1} is strictly increasing in y and there exists y_l, y_h such that $\pi^{L1}(y_l) < 0$ and $\pi^{L1}(y_h) > 0$. If there exists a x such that $\pi^{L1}(x) = 0$ then define κ such that $\pi^{L1}(\kappa) = 0$. If not, (may happen since Q is not continuous), let κ be such that $\pi^{L1}(\kappa) < 0$ and $\pi^{L1}(\kappa + \epsilon) > 0$ for any $\epsilon > 0$. Such a κ is unique as π^{L1} is strictly increasing. Thus, L1 types play according to a threshold strategy with cutoff κ .

Assume true for k-1.

Consider k.

Lk types expect everyone to attack if $y > \kappa$.

If $y < \kappa$, payoff to attacking is $y \cdot 0 - t < 0$

If $y \geq \kappa$ then payoff to attacking is $y \cdot 1 - t > 0$

The result follows by induction.