

Rationally Inattentive Acquisition of Information in Dynamic Global Games

By

Examination Number : 32383

A Dissertation submitted to,
the Department of Mathematics at
London School of Economics and Political Science, United Kingdom
for the degree of Master of Science

August 2024

Supervisor(s) : Prof. Olivier Gossner

ABSTRACT

This dissertation integrates the concept of rational inattention into the framework of dynamic global games to examine the role of information in financial crises and the implications of information acquisition costs. By allowing agents to optimally allocate limited attention to process relevant information, this work extends the traditional models of dynamic global games pioneered by Angeletos, Hellwig, and Pavan (2007).

Dynamic global games provide a comprehensive framework for analyzing strategic interactions among agents over time, where decisions are based on private information and the observed actions of others, particularly under uncertainty. Unlike static games, which capture interactions at a single point, dynamic games account for the evolution of strategies and information flows across multiple periods. This approach is essential for understanding phenomena like currency attacks, bank runs, and financial crises, where the timing and sequence of actions are critical.

By incorporating the dynamics of information acquisition costs into this framework, our research more accurately mirrors real-world scenarios where obtaining and processing information incurs significant costs. Emphasizing on the temporal nature of strategic decision-making, our theoretical framework aims to reduce strategic uncertainty and promote a unique equilibrium, providing a more accurate depiction of market dynamics.

This work advances the theoretical literature on dynamic global games, offering deeper insights into strategic decision-making in financial markets and highlighting the importance of studying dynamic games over static frameworks.

Contents

Page

1 Introduction **1**

1.1 Motivation 1

1.2 Reasoning Multiplicity 3

1.3 Uniqueness vs Multiplicity 3

1.4 Costly Information Acquisition 4

2 Coordination Games **6**

2.1 A Simple Coordination Game Problem 7

2.2 Intuition behind Global Games 8

2.3 Binary Action Global Game Model for Two Players 9

2.3.1 Baseline Model Setup 9

2.3.2 Defining Bayesian Nash Equilibrium for G(θ) 10

2.3.3 Examining the nature of perturbations 13

3 Applied Macroeconomic Model for a Continuum of Players **14**

3.1 Static Framework 14

3.2 The Benchmark Case of Complete Information 15

3.3 Introducing Morris- Shin's strategic uncertainty 16

3.4 Important Concept of Equilibrium 17

3.4.1 Describing Monotone Equilibria 18

3.5 Applications to Macroeconomics 22

4 Comprehending Dynamic Global Games **23**

4.1 Setting up the Dynamic Framework 23

4.2 Reasonable conditions for Optimality 25

4.2.1 Implications of the Lemma 28

4.2.2 Equilibrium with Multiple attacks 29

4.3 Alternate Expressions of Equilibrium 30

4.4 Embracing Monotonicity 32

4.4.1 Interpretation of Lemma 34

4.4.2 An Instance of Multiplicity 36

5 Rationally Inattentive Acquisition of Information **38**

5.1 A Non Technical Introduction to Rational Inattention 38

5.2 Foundations - The Rational Inattention Game 39

5.2.1 Curious Case of Flexible Information Structure 39

5.2.2 Conditional Choice probabilities as a response to Information Signals 40

5.3 Cognitive Costs of Information 41

5.4 Objectives of the RI Game 42

5.5 Amending Bayesian Correlated Equilibrium 43

5.6 The Case of Shannon's Entropy Reduction Functional 43

5.7 Plausible Extensions in the Dynamic Framework 44

5.7.1 Setup 45

5.7.2 Dynamic Strategies 46

5.8 Reaching Equilibrium Uniqueness 47

6 Conclusion **50**

References **52**

Chapter 1

Introduction

1.1 Motivation

Imagine a crowded theatre where someone yells " Fire !" Panic ensues, and people rush to the nearby exits, sometimes causing more harm than the fire itself. This scenario might come across as a metaphor for many financial and political crises we witness in the world today. Crises like bank runs, where fears of insolvency prompt a sudden mass withdrawal of deposits, or currency attacks, where speculators aggressively bet against a nation's currency, or even political revolts, where groups rise up against a dictator, can wreak havoc on economies and create widespread turmoil. These fascinating real-world economic phenomena often lack clear-cut explanations and can be effectively modeled using coordination games with incomplete information. In such games, the payoffs for each participant hinge not only on their own actions and those of others but also on their beliefs about key economic fundamentals. These beliefs extend beyond just the basics—they include layers of higher-order beliefs, such as what players think about the fundamentals, what they believe others think about the fundamentals, and even what they believe others believe about others' beliefs, creating a complex web of expectations and strategies.

Economic fundamentals can be thought of as the underlying factors that influence the structure of the game and the rewards at stake. For instance, in the context of a bank run, these fundamentals might include the bank's ability to fulfill its obligations without external intervention. Drawing on the insights from Mertens and Zamir (1985), we can describe a

player's type in a game with incomplete information through an infinite hierarchy of beliefs for all possible states of the world.

However, as we delve into the reasons behind such dramatic collective shifts in strategic behavior and seek to find optimal strategies within this infinite space of beliefs, the complexity of the analysis can quickly become overwhelming and nearly impossible to manage. Therefore, it is extremely valuable to identify strategic environments that capture the essence of these higher-order beliefs in economic settings while remaining simple enough to allow for manageable analysis. This is where global games come into play.

First introduced by Carlsson and Van Damme (1993), global games have proven to be a powerful tool for modeling problems in coordination game theory, such as currency crises (Morris and Shin, 1998), bank runs (Goldstein and Pauzner, 2005), foreign direct investment decisions (Dasgupta, 2007), and even regime changes and political revolts (Angeletos, Hellwig, and Pavan, 2007). In these games, agents' payoffs depend both on the state of the economy—viewed as economic fundamentals—and on the actions of others. A key assumption in these macroeconomic models is that players are unaware of the actual realization of the economic fundamental, making it crucial for them to gather more information before making decisions. To do this, they rely on public signals (like news reports, public announcements, and discussion forums) and private information sources, which provide only a noisy estimate of the true state of the economy.

As players receive these noisy signals, they not only form their own strategies but also make inferences based on the perceived beliefs of their opponents. This particular alteration in the information structure of the game generates a rich array of higher-order beliefs, guiding players toward a unique equilibrium. This prediction of a unique equilibrium is a striking departure from the expectations in coordination games with complete information, where multiple equilibrium outcomes are possible. This uniqueness makes global games especially useful for developing policy recommendations with real-world impact, offering a clear and actionable framework for addressing complex economic issues.

1.2 Reasoning Multiplicity

Morris and Shin (2000) delve into the foundations of the multiple equilibria literature in macroeconomics, focusing on game theory. They highlight two key assumptions that often lead to multiple equilibria: strategic complementarities, where players benefit from coordinating their actions, and the notion of common knowledge about economic fundamentals. Traditional rational expectations theory, including Nash Equilibrium, assumes that players have a clear understanding of each other's actions, allowing them to align their strategies and beliefs, which in turn creates multiple possible equilibria. This multiplicity arises from the uncertainty in beliefs, as different sets of beliefs can support various equilibrium outcomes. The concept of "sunspots" is used to illustrate how self-fulfilling beliefs can influence strategies, even if these beliefs are not directly tied to economic fundamentals.
However, the "sunspots" models do not explain how these beliefs form or why they shift, nor do they provide a strong link between economic sentiment and the actual strength of fundamentals. Morris and Shin (2000) question the common knowledge assumption and introduce noisy private signals to introduce uncertainty about others' actions. This uncertainty disrupts perfect coordination, leading to a unique equilibrium when private information is precise enough. This alternative approach, which reduces the reliance on common knowledge and emphasizes the role of information asymmetry, will be explored further in the subsequent sections of this dissertation.

1.3 Uniqueness vs Multiplicity

In modern economic theory, particularly within the study of global games, the concept of a stable, unique equilibrium is pivotal for understanding how agents' expectations and actions converge. A unique equilibrium provides a clear, unambiguous prediction of outcomes, ensuring that players can consistently coordinate their strategies despite observing noisy

signals about the fundamentals and facing high chances for mis-coordination. Moreover, this uniqueness implies robustness, making the equilibrium outcome less sensitive to small perturbations in the underlying parameters or information structure, which is essential for reliable application in real-world scenarios.

From a policy perspective, a unique equilibrium significantly simplifies the design and implementation of interventions. Policymakers can be more confident in the predicted outcomes of their actions in environments that tend toward a unique equilibrium. In contrast, multiple equilibria introduce strategic uncertainty, making the impact of policy measures highly unpredictable and potentially hindering effective decision-making. Additionally, the presence of multiple equilibria complicates empirical analysis, as identifying the specific equilibrium being played in the real world becomes challenging, thereby reducing the explanatory power of the model.

1.4 Costly Information Acquisition

Recent developments in global games have expanded beyond the original framework, particularly regarding costly information acquisition in dynamic settings. Szkup and Trevino (2015) pioneered research in this field by integrating information costs into static global games of regime change with endogenous information evolution. Their model found a unique, risk-dominant, but inefficient symmetric equilibrium under mild conditions, though it was restricted by its additive information structures.

Building on the theory of Rational Inattention (Sims, 2003), which limits information processing to a feasible reduction in entropy, a more flexible framework for information costs has emerged. This framework is studied in linear-quadratic (Maćkowiak and Wiederholt, 2009) and binary-action forms (Matějka and McKay, 2015, Hébert and Woodford, 2023). Our model, discussed in Chapter 5, will extend the latter approach to include cognitive costs over multiple periods.

Yang (2015) challenged the integration of Rational Inattention into coordination games, showing that using entropy reduction as a cost functional can reverse uniqueness and selection results in global games. However, the obtained multiplicity results were not entirely unanticipated, since Yang (2015) had worked out a form of perturbation of common knowledge which was very distinct from what was proposed in Carlsson and Van Damme (1993). They argue that specific to the benchmark models of Global Games, when information acquisition is rigid and players are compelled to assign equal attention to all possible realizations of the fundamental despite some values being vague and preposterous, players coordinate only in choosing overall precisions. They cannot coordinate to achieve approximate common knowledge for any "cutoff event" and hence a unique equilibrium persists. This contrasts sharply with their flexible information structure which allows players to acquire more information concerning a certain realization if they wish to do so.

This dissertation, however, adopts a different approach: we incorporate costly information acquisition in dynamic global games, extending the standard perturbations of Carlsson and Van Damme (1993) across multiple periods and use a non-parametric cost functional, inspired by Hoshino (2018). While Angeletos, Hellwig, and Pavan (2007) explored dynamic global games without considering information costs, their model led to multiple equilibria, which is less desirable as discussed in Section 1.3. We argue that introducing information costs can lead to unique equilibrium outcomes by reducing coordination among players, an area still underexplored and promising for future research.

The dissertation is structured as follows: Chapter 2 introduces Coordination Games and explains the motivation behind studying Global Games. Chapter 3 reviews the static global game framework with a macroeconomic model of Regime Change. Chapter 4 explores the dynamic game setup of Angeletos, Hellwig, and Pavan (2006) and monotone equilibria. Chapter 5 extends this model with costly Rational Information Acquisition and assesses its impact. Chapter 6 concludes.

Chapter 2

Coordination Games

On numerous occasions, we choose to participate in certain activities simply because they are popular. The Scottish philosopher David Hume referred to this behavior as following conventions (Latsis, 2009). For Hume and many philosophers after him, language was a prime example. We call an "apple" an apple and an "orange" an orange because everyone else does. If the majority switched and started calling an apple an orange and vice versa, we would likely follow suit. This is the basic payoff of coordination.

In contemporary game theory, such scenarios are classified as coordination games, where multiple stable conventions, or equilibria, can exist. These coordination problems are common in significant social situations, and the importance of aligning incentives, known as strategic complementarities, has been highlighted in various real-world crises. For instance, when the US invaded Grenada in the 1980s, communication issues arose because the marines and army used incompatible radios, leading to inefficient and cumbersome coordination. Similarly, during the global financial crisis, shadow banking lenders reduced their lending when others did, amplifying the market freeze. European sovereign debt holders also reduced their holdings when others did, driving up interest rates and increasing the risk of default.

These crises illustrate how lack of complete information and uncertainty about market conditions can lead to disastrous outcomes. If participants had perfect knowledge of the situation and could accurately predict others' actions, these coordination failures might have been avoided. With this understanding, let's explore the equilibrium dynamics in a simple coordination game where complete knowledge is assumed.

2.1 A Simple Coordination Game Problem

To illustrate the setup for the model, let's examine a payoff matrix for a basic coordination game involving two players, A and B, who must independently choose between two actions: Up (U) and Down (D) for Player A; Left(L) and Right(R) for Player B. The payoff matrix is as follows:

	L	R
U	(0,0)	(0,-1)
D	(-1,0)	(1,1)

In comparative statics, two outcomes emerge in a coordination game: (U,L) and (D,R), where no player has an incentive to deviate. Both are strict Nash equilibria, with (D,R) being Pareto efficient, meaning no player can improve their payoff without reducing their opponent's, while (U,L) is risk-dominant, as it involves less risk given uncertainty about the other player's actions.

A key characteristic of this simple coordination game is the crucial role of confidence and expectations in shaping players' choices. When structural changes affect one player's payoff, a natural transmission process occurs, leading to consistent reactions from all players and resulting in positive correlations. Additionally, if we assume a preference ranking—such as R $\succ$ L for player A, and D $\succ$ U for player B—these games exhibit strategic complementarities, where a "dominant" choice by one player increases the marginal payoff for the other to also choose their dominant action. They also show positive spillovers, where an increase in one player's choice enhances the other's payoff.

Recognizing the importance of confidence and expectations in determining rational behavior, and understanding that games with complete information are the limit of games with incomplete information, Carlsson and Van Damme (1993) argued that introducing private information into this system could impact the multiplicity of equilibria.

2.2 Intuition behind Global Games

A coordination game with complete information typically assumes that all players have common knowledge about the state and the payoffs at stake. However, instead of assuming players are always perfectly informed and rational about their choices, we can introduce a small amount of continuous noise to the payoffs. This perturbation leads us to consider an ex-ante stage where payoffs are observable before players receive their signals. Such games are known as Global Games (Carlsson and Van Damme, 1993). Unlike coordination games with perfect common knowledge, global games allow theorists to move beyond the assumption of perfect coordination of actions and beliefs, incorporating higher-order beliefs into strategic settings.

To this extent, let us finally consider the parametric class of games, G(θ) with two players A and B, having the following structure : (Carlsson and Van Damme, 1993)

	Invest	Not Invest
Invest	(θ,θ)	(θ-1,0)
Not Invest	(0,θ-1)	(0,0)

In this example, it's evident that the payoffs depend on θ representing an economic fundamental, such as the current state of the financial market. The players have two possible actions: Invest and Not Invest. Key features of this coordination game with a fixed fundamental are as follows:

- When $\theta < 0$, (Not Invest, Not Invest) is the unique dominant strategy.
- When $\theta > 1$, then (Invest, Invest) is the unique dominant strategy.
- When $\theta \in [0, 1]$, there are two Nash Equilibria that can be Pareto ranked; Invest is risk dominant if $\theta \geq \frac{1}{2}$, Not Invest is risk dominant otherwise if $\theta \leq \frac{1}{2}$

This parameterized class of games provides the foundation to introduce global games.

2.3 Binary Action Global Game Model for Two Players

Morris and Shin (2002) define global games as games of incomplete information where the type space is generated by the players each observing a noisy private signal of the underlying state. Before we establish the baseline two player two action model of a global game, we must point out that while we do not assume common knowledge amongst the players in terms of the economic fundamentals, we do assume common knowledge in terms of how these fundamentals as well as the idiosyncratic noise technology is randomly distributed.

2.3.1 Baseline Model Setup

The following setup is due to Carlsson and Van Damme (1993). Consider the same parametric class of games G(θ) with two players, call them Player 1 and 2, (as discussed in the section 2.2) :

	Invest	Not Invest
Invest	(θ,θ)	(θ-1,0)
Not Invest	(0,θ-1)	(0,0)

Unlike the hypothesis put in place in the Section 2.2, we now assume no common knowledge amongst the players in terms of θ, as they only acquire a private signal $x_i = \theta + \sigma\epsilon_i$. Here, ϵ_i, or the idiosyncratic noise(while observing θ) is identically and independently normally distributed $\epsilon_i \sim N(0,1)$. θ is assumed to be uniformly distributed on the real line which, although not a standard assumption, presents no technical difficulties whatsoever. Therefore, the player $i's$ private signal x_i is informative on the underlying state of nature θ and on the other players' signal x_{-i} since these signals correlate with θ. In particular, we can imply from standard statistical theory that, the expected value of the fundamental to a player i

contingent upon the acquisition of a private signal is

$$E[\theta|x_i] = x_i \tag{2.1}$$

The player $i's$ posterior belief about the fundamental is given by :

$$\theta|x_i \sim N(x_i, \sigma^2) \tag{2.2}$$

and finally, the player's beliefs about opponents' signal conditional on own signal is given by

$$x_{-i}|x_i \sim N(x_i, 2\sigma^2) \tag{2.3}$$

2.3.2 Defining Bayesian Nash Equilibrium for G(θ)

In this simple binary action global game, Carlsson and Van Damme (1993) take the pure strategy for a player i to be a mapping $s_i : \mathbb{R} \rightarrow$\{Invest, Not Invest\}. Then we suppose player i believes that their opponents, $-i$, are playing threshold investment strategies with cutoff k such that

$$s^k_{-i}(x_{-i}) = \begin{cases} \text{Invest} & x_{-i} > k \\ \text{Not Invest} & x_{-i} \leq k \end{cases} \tag{2.4}$$

Clearly, now that we have already established from 2.3 and 2.4 that, $x_{-i}|x_i \sim N(x_i, 2\sigma^2)$ and that player i believes their opponents will not invest when their private signals $x_{-i} \leq k$, then by the properties of the standard normal variate, player i would believe that player $-i$, will not invest with probability

$$\phi(\frac{1}{\sqrt{2}\sigma}(k - x_i)) \tag{2.5}$$

which further implies that player $i's$ expected payoff if they decides to invest is given by

$$x_i[1 - \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i))] + (x_i - 1)\phi(\frac{1}{\sqrt{2}\sigma}(k - x_i))] = (x_i - \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i)) \tag{2.6}$$

so that $i's$ best response to such an opponent strategy $s^k_{-i}(x_{-i})$ is

$$s^{br}_i(s^k_{-i}(x_{-i})|x_i) = \begin{cases} \text{Invest} & x_i > \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i) \\ \text{Not Invest} & x_i \leq \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i) \end{cases} \tag{2.7}$$

Note that the Equation 2.7 simply reiterates the argument that player i would like to invest if their payoff from investing is greater than 0, and would refrain from doing so, if this payoff comes out to be ≤ 0. Hence, the equation $x_i - \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i) = 0$ represents the indifference expression for any player i, which also by extension has a unique solution in x_i, $b(k)$. Hence implicitly, the best response of player i is to follow a cut-off (or later, referred to as a monotone) strategy with threshold equal to $b(k)$:

$$s_i^{br}(s_{-i}^k(x_{-i})|x_i) = \begin{cases} \text{Invest} & x_i > b(k) \\ \text{Not Invest} & x_i \leq b(k) \end{cases} \tag{2.8}$$

Now, following directly from this Indifference Criterion 2.7, and the properties of a standard normal variate ϕ, we can infer that :

1. When $k \to -\infty \implies \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i) \to 0 \implies b(k) \to 0$
   This reflects the situation that if one player is choosing to always invest, the other player will only invest if their signal x_i is positive.

2. On the other hand, when $k \to \infty \implies \phi(\frac{1}{\sqrt{2}\sigma}(k - x_i) \to 1 \implies b(k) \to 1$
   This reflects the situation that if one player is choosing to always not invest, the other player will only invest if their signal x_i is greater than 1.

3. Finally, when $k = \frac{1}{2}$, it only implies that $b(k) = \frac{1}{2}$. Now, this reflects the case that when player acquires a private signal $x_i = \frac{1}{2}$, then $\mathbb{E}[\theta|x_i] = \frac{1}{2}$, which implies that according to $i's$ beliefs, his opponents, parameterized by $-i$, will not invest with probability $\frac{1}{2}$.

Additionally, by total implicit differentiation, we get :

$$b'(k) = \frac{1}{1 + \frac{\sqrt{2\Pi}}{\phi(\frac{1}{\sqrt{2}\sigma}(k - x_i)}} \in (0, 1)$$

Clearly, from the analysis done above, it is not difficult for us to conclude that the Bayesian Nash Equilibrium is unique for both players choosing to invest only when they observe private signals greater than $\frac{1}{2}$. Moreover, from the conclusions put forth by Milgrom and Roberts,

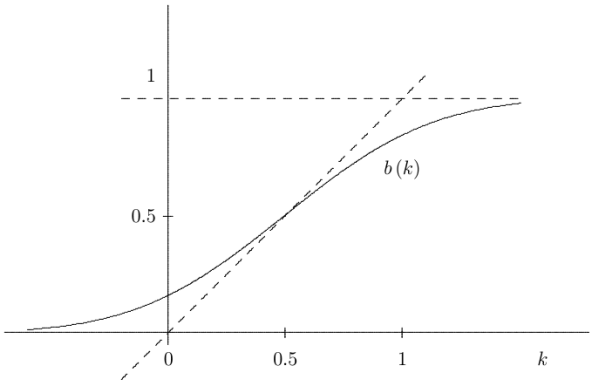

Figure 2.1: Function $b(k)$, Morris and Shin(2003)

1990, the strategy with threshold $1/2$ is in fact the unique strategy surviving iterated deletion of (interim) strictly dominated strategies. since a strategy s_i survives n rounds of iterated deletion of strict dominated strategies if and only if

$$s_i(x_i) = \begin{cases} \text{Invest} & x_i > b^n(1) \\ \text{Not Invest} & x_i \leq b^n(0) \end{cases} \tag{2.9}$$

where, $b^n(k) = b(b(\dots b(k)))$.

The proof can be done using the arguments by Milgrom and Roberts (1990) or by the following Induction argument inspired from Morris and Shin (2003): It is trivial to prove the claim for $n = 1$, since "Not Invest" is a dominant strategy for a player for the expected value of θ less than 0. Now, suppose the claim is true for arbitrary n. If a player knew that their opponent would choose action "Not Invest" if they had observed a signal less than $b^{n-1}(1)$, their best response would always be to choose action "Not Invest" if their signal was less than $b(b^{n-1}(1))$ Because b$(\cdot)$ is strictly increasing and has a unique fixed point at $\frac{1}{2}$, $b^n(0)$ and $b^n(1)$ both tend to $\frac{1}{2}$ as $n \to \infty$.

With a uniform prior, each player always believes there is a 50% chance their opponent has a higher signal and a 50% chance they have a lower signal. This belief remains constant regardless of the noise level but shifts to certainty when noise approaches zero: without noise, players are certain their opponent will choose the same action. This finding highlights that even small variances can significantly impact decisions due to the potential presence

of dominant strategies. Consequently, the probability that either player invests is $\phi(\frac{\frac{1}{2}-\theta}{\sigma})$. Investment decisions are independent given θ.

2.3.3 Examining the nature of perturbations

The two-player, two-action coordination games we discuss are called "Global Games" because they work with various payoff profiles as long as idiosyncratic noise is consistently introduced. In Carlsson and Van Damme (1993), this approach led to two main results: Firstly, a unique action profile survives the iterated deletion of strictly dominated strategies, and Secondly, this profile remains unaffected by idiosyncratic noise in the limit.

Carlsson and Van Damme (1993) relaxed the common knowledge assumption about economic fundamentals (parameterized by θ), assuming players had common prior beliefs and received conditionally independent private signals. As noise decreased, the perturbed types approximated common knowledge types, similar to the universal type space.

As noise approaches zero, the threshold strategy equilibrium becomes sensitive to higher-order beliefs, akin to Rubinstein (1989). This raises the question of when an action profile, which is an equilibrium under common knowledge of θ, uniquely survives in perturbed games after iterated deletion of strictly dominated strategies.

Yildiz and Weinsten (2004) demonstrated that any action profile surviving the iterated deletion in the complete information game meets this criterion. Their work emphasizes that unique equilibrium selection in global games results from a specific relaxation of the common knowledge assumption and highlights the stability of higher-order beliefs.

We now move to a macroeconomic model of Regime Change with a continuum of players and two actions, offering a broader view of global games and information's role in strategic decision-making.

Chapter 3

Applied Macroeconomic Model for a Continuum of Players

3.1 Static Framework

Hypothesize two possible regimes - the status quo, and the alternative. Let there be a measure - one continuum of players in the model, indexed by $i \in [0, 1]$, so that each player must either choose to reject the current regime or support it. We label these plausible actions as "Attack (A)" or "Not Attack (NA)" respectively. We assume all agents to choose an action simultaneously, in that the game is static.

Further, the set of possible outcomes is given by $R \in \{0, 1\}$, where $R = 0$ represents the case that the status quo is retained, and $R = 1$ represents the alternate case that it is abandoned. Likewise, the set of possible actions for a player i is given by $a_i \in \{0, 1\}$, where $a_i = 0$, represents "Not Attack" and $a_i = 1$, represents "Attack".

Normalize the payoffs to player for not attacking to be 0. On the contrary, the payoffs to the player from attacking are taken to be $1 - c > 0$, if the status quo is abandoned (R = 1), and $-c < 0$, if it is sustained (R = 0) with c characterizing the relative cost of attacking. In the context of a political revolt game, one may think of such a relative cost as the amount of materialistic resources needed to pull off the attack. In concise, the utility of player i is

$$u_i = a_i(R - c)$$

Intuitively, the proportion of players attacking the status quo would then be given by $A = \int a_i di \in [0,1]$; the strength of the status quo regime parameterized by $\theta \in \mathbb{R}$. So that the status quo is abandoned if and only if $A \geq \theta$. This only reconciles the argument that if the mass of players attacking the status quo is higher than or even equal to the strength of the regime, the status quo will cease to exist. Additionally, it is not wrong to conclude that a player would have a strong incentive to attack the status quo, if and only if it eventually collapses, and furthermore, the status quo will only collapse if and only if a sufficient proportion of players coordinate their actions and attack together. Hence, the players' actions are strategic complements.

To summarise, we can rewrite the expression for the payoff function to an individual player $i \in [0,1]$ as :

$$u_i = U(a_i, A, \theta) = \begin{cases} a_i(1-c) & A \geq \theta \\ -a_i c & A \leq \theta \end{cases} \tag{3.1}$$

Precisely, $U(a_i, A, \theta)$ represents the utility to player i for choosing an action a_i contingent upon the mass of the attack A, and the strength of the regime θ. Moreover, the best response for any player $i \in [0,1]$ is simply given by

$$g(A, \theta) = arg \max_{a_i \in \{0,1\}} U(a_i, A, \theta) = \begin{cases} 1 & if A \geq \theta \\ 0 & if A \leq \theta \end{cases}$$

Owing to the existence of strategic complementarities, both U and g are strictly increasing in A, mass of attack.

3.2 The Benchmark Case of Complete Information

If all players have complete knowledge of the regime's strength, represented by θ, then θ is common knowledge and taken as given. Several scenarios arise depending on θ's value. When $\theta \leq \underline{\theta}$ (where $\underline{\theta} \equiv 0$), the regime is so weak that every player has an incentive to attack, leading to the inevitable collapse of the status quo. The unique equilibrium here is

that all players attack. On the other hand, when $\theta \geq \bar{\theta}$ (where $\bar{\theta} \equiv 1$), the regime is so strong that attacks are futile, and the equilibrium outcome is that no player attacks.

However, when θ takes intermediate values between $\underline{\theta}$ and $\bar{\theta}$, i.e., $\theta \in [\underline{\theta}, \bar{\theta}]$, the regime, while strong, remains vulnerable to a large enough attack. This leads to multiple equilibria driven by self-fulfilling expectations. In one scenario, players believe everyone else will attack, prompting them to do the same, causing the regime to collapse. In another, players expect no one to attack, so they refrain, preserving the status quo.

Morris and Shin (2003) characterize the range $(\underline{\theta}, \bar{\theta}]$ as the set of "critical fundamentals", where multiple equilibria occur under common knowledge. Each equilibrium is sustained by different self-fulfilling expectations about others' actions. When θ is common knowledge, players can perfectly predict each other's actions, allowing for coordination on multiple equilibria.

However, in most real-world scenarios modeled as game-theoretic stage games, players seldom have perfect knowledge of the economic fundamentals they face. Thus, assuming players have no need to gather information about the regime's strength is an unrealistic starting point.

3.3 Introducing Morris- Shin's strategic uncertainty

Building on the literature initiated by Morris and Shin (1998) regarding the uniqueness of self-fulfilling currency crises, we now focus on the scenario where θ is not common knowledge and only incomplete information is available. In this case, players are uncertain about the strength of the status quo and their opponents' actions, so they rely on heterogeneous private noisy beliefs about θ to reduce uncertainty and make strategic decisions between attacking or not attacking. Mathematically, this can be expressed as follows:

Players initially draw a common prior belief about the strength θ of the status quo, hypothesized to be uniformly distributed over the real line. Each player $i \in [0, 1]$ observes a private

signal:

$$x_i = \theta + \zeta_i$$

where ζ_i represents private noise in the estimation of θ. This noise is normally distributed, $\zeta_i \sim N(0, \frac{1}{\beta})$, and is independent of θ across all players. Thus, the private signal x_i serves as a sufficient statistic for the private information each player acquires.

With a continuum of players, the collective information $(x_i)_{i\in[0,1]}$ across the economy is sufficient to determine θ, though this knowledge is uniformly spread. Additionally, we assume that players may also have access to public information, such as news or public announcements. Each player observes a public signal:

$$z = \theta + \mu$$

where $\mu \sim N(0, \frac{1}{\alpha})$. where $\mu \sim N(0, \frac{1}{\alpha})$. The public signal z represents the realization of the public signal for each player and defines their initial prior before acquiring private signals. The noise μ in the public signal is normally distributed and independent of both θ and the private noise ζ. The information structure can thus be parameterized by the precisions (β, α), given the normality assumptions.

3.4 Important Concept of Equilibrium

Definition 3.4.1 (*Symmetric Bayesian Nash Equilibrium*)**.** *The equilibrium, in this context, is defined by a strategy $a(\cdot)$ and an aggregate mass of attack $A(\cdot)$ such that:*

$$a(x, z) \in arg\max_a \mathbb{E}[U(a, A(\theta, z), \theta)|x, z] \tag{3.2}$$

Furthermore,the equilibrium size of attack is given by

$$A(\theta, z) = \int a(x, z)\sqrt{\beta}\phi(\sqrt{\beta}[x - \theta])dx \tag{3.3}$$

This definition simply ensures that any action strategy will only define Nash equilibrium if and only it maximizes the expected utility to a player, conditional on the acquisition of information, both private and public.

3.4.1 Describing Monotone Equilibria

In this section, we continue the discussion on the static global game model of Regime Change by characterizing the features of monotone equilibria associated with the game.

In a monotone equilibrium, for any realization of the public signal, z, there exists a threshold $x^*(z)$ such that a player would only have an incentive to attack the status quo if $x \leq x^*(z)$. Think of it in a way that only when a player's private belief about the strength of the status quo is not sufficiently high, they will find it profitable to attack. Additionally, the marginal size of the attack, decreases as the strength θ increases, so that there is a threshold $\theta^*(z)$ such that the status quo is abandoned if and only if $\theta \leq \theta^*(z)$. Furthermore, such an equilibrium can be classified using the private threshold x^* and the state threshold, θ^*.

Hence, in order to completely describe this equilibrium, we proceed in three stages, as inspired by Angeletos and Werning (2006). In the first stage, we classify the state threshold θ^* for a given x^*; Following which, in the next stage, we classify the private threshold x^* for a given θ^*. And lastly, in the final stage, we use the previous results and achieve proving equilibrium uniqueness and, by extension, its existence.

Hence, consider the first stage. Given some realization of the fundamental θ, signifying the strength of the regime, as well as the public signal z, the aggregate size of the attack is essentially the measure of players choosing to attack, which is further nothing but the measure of players receiving signals $x \leq x^*(z)$. Then by properties of the standard normal variate, we can write:

$$A(\theta, z) = \phi(\sqrt{\beta(x^*(z) - \theta)}) \tag{3.4}$$

where β is the precision of the private information. We must note that $A(\theta, z)$ is decreasing in θ, so that the regime change will occur if and only if $\theta \leq \theta^*(z)$ (strength of the regime being sufficiently low) is the unique solution to

$$A(\theta^*(z), z) = \theta^*(z)$$

which further simplifies to

$$x^*(z) = \theta^*(z) + \frac{1}{\sqrt{\beta}}\phi^{-1}(\theta^*(z))$$

As a next stage, we turn to examining equilibrium x^* for a given θ^*.
Having known that the regime change occurs if and only if $\theta \leq \theta^*(z)$, the payoff to a marginal player then follows from the functional we described through Equation 3.1

$$\mathbb{E}[U(a, A(\theta, z), \theta)|x, z] = a(Prob[\theta \leq \theta^*(z)|x, z] - c) \tag{3.5}$$

Remember that β and α denote the precision of private and public information received by the player, respectively. Clearly, the model so far carefully illustrates the nature of the signals which a player can acquire by participating in the game. Once he acquires these signals, he uses them to update his beliefs about the strength of the status quo, and subsequently form limit-informed decisions. Such an updated knowledge about the strength of the regime takes the form of the posterior belief of the player, given by

$$\theta|x, z \sim N(\delta x + (1 - \delta)z, \alpha^{-1}) \tag{3.6}$$

where $\delta \equiv \frac{\beta}{\beta+\alpha}$ represent the relative precision of private information, and $\gamma \equiv \beta + \alpha$ is the overall precision of information. Intuitively, the posterior (updated) probability of regime change would also be nothing but the probability that is conditional on the private and public signal received by each player. Hence, we have:

$$\mathbb{P}[\theta \leq \theta^*(z)|x, z] = 1 - \phi(\sqrt{\gamma}(\delta x + (1 - \delta)z - \theta^*(z))) \tag{3.7}$$

which is clearly monotone with respect to the strength of the private signal that a player acquires.

Particularly, a player would only attack if and only if $x \leq x^*(z)$ with $x^*(z)$ solving the indifference criterion.

$$b \cdot \mathbb{P}[\theta \leq \theta^*(z)|x^*(z), z] - c = 0 \tag{3.8}$$

It must be noted here that the Indifference Criterion is reflective of the situation that a player is indifferent between attacking and not attacking, i.e, when the payoff to the player is precisely 0.

Finally, as we substitute the expression of the posterior probability from Equation 3.7 , as well as the definitions of relative and overall precisions of private signals mentioned earlier into the first form of the Indifference criterion 3.8, we obtain :

$$\phi\left(\sqrt{\beta+\alpha}\left(\frac{\beta}{\beta+\alpha}x^*(z) + \frac{\alpha}{\beta+\alpha}z - \theta^*(z)\right)\right) = \frac{b-c}{b} \tag{3.9}$$

Final Stage : Combine Equations 3.7 and 3.8, to conclude that $\theta^*(z)$ can be sustained in equilibrium if and only if it solves

$$G(\theta^*(z), z) = g$$

where

$$g = \sqrt{1+\frac{\alpha}{\beta}}\phi^{-1}(\frac{1-c}{b}) \tag{3.10}$$

and

$$G(\theta, z) \equiv \frac{\alpha}{\sqrt{\beta}}(z-\theta) + \Phi^{-1}(\theta) \tag{3.11}$$

With $\theta^*(z)$ given by 3.9, $x^*(z)$ is then given by 3.7.

Illustration of Existence and Uniqueness

Firmly we can now establish the existence and uniqueness of the equilibrium by taking into account the properties of the function G. From the expression for G iterated in Equation 3.11, we deduce that, for every $z \in \mathbb{R}$, $G(\theta, z)$ is continuous in θ, with $G(\underline{\theta}, z) = -\infty$ and $G(\bar{\theta}, z) = \infty$, which ensures the existence of a particular solution due to the Intermediate

Value Theorem, and any such solution would satisfy $\theta^*(z) \in (\underline{\theta}, \bar{\theta})$. Having proved existence, we now turn to uniqueness.

Consider the Equation 3.11 and take the partial derivative of G with respect to θ, so that we obtain

$$\frac{\partial G(\theta, z)}{\partial \theta} = \frac{1}{\phi(\Phi^{-1}(\theta))} - \frac{\alpha}{\sqrt{\beta}} \tag{3.12}$$

Since $\max_{w \in \mathbb{R}} \phi(w) = \frac{1}{\sqrt{2\pi}}$, then if $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$, it is directly implied that $\frac{\partial G}{\partial \theta} \geq 0$, which in other words, mean that G is strictly increasing in θ. Hence, this ensures a unique solution to the expression obtained above.

However, if it so happens that $\frac{\alpha}{\sqrt{\beta}} > \sqrt{2\pi}$ G is non-monotonic in θ and there is an interval $(\underline{z}, \bar{z})$ such that Equation 3.7 admits multiple solutions $\theta^*(z)$ whenever $z \in (\underline{z}, \bar{z})$, or a unique equilibrium outcome otherwise. In conclusion, we find that the monotone equilibrium is unique if and only if $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$

More formally, these results have been summarized as follows :

Proposition 3.4.1 (Morris and Shin, 1998)**.** *Let $\sqrt{\frac{1}{\beta}}$ and $\sqrt{\frac{1}{\alpha}}$ denote the standard deviations of the private and the public signals, respectively. Then, there always exist a monotone equilibrium and it is unique if and only if $\frac{\alpha}{\sqrt{\beta}} \leq \sqrt{2\pi}$.*

Additionally, if we take the limits as standard deviation in private signal $\sqrt{\frac{1}{\beta}} \to 0$ for a given $\sqrt{\frac{1}{\alpha}}$, or $\sqrt{\frac{1}{\alpha}} \to \infty$ for a given $\sqrt{\frac{1}{\beta}}$, we get $\frac{\alpha}{\sqrt{\beta}} \to 0$ and $\sqrt{(\beta + \alpha)/\beta} \to 1$ in either cases. Now, the condition 3.9 discussed earlier then implies that $\theta*(z) \to \hat{\theta} = 1 - c/b$ for all z. All of this put together allows us to further establish the following result, which we refer to as the Morris and Shin Limit :

Proposition 3.4.2 (Morris and Shin Limit)**.** *In the limit as either $\sqrt{\frac{1}{\beta}} \to 0$ for a given $\sqrt{\frac{1}{\alpha}}$, or $\alpha \to \infty$ for a given $\sqrt{\frac{1}{\beta}}$, there is a unique monotone equilibrium, in which the regime changes if and only if $\theta \leq \hat{\theta}$, where $\hat{\theta} = 1 - \frac{c}{b} \in (\underline{\theta}, \bar{\theta})$ Moreover, from the Iterated Elimination of Strictly Dominated Strategies, we know that when the monotone equilibrium*

is unique, there is no other equilibrium.

3.5 Applications to Macroeconomics

The simplified model of regime changes helps us grasp the role of coordination dynamics across various real-world scenarios, even if the game structure is somewhat stylized. In self-fulfilling currency crises models (Obstfeld, 1986, Obstfeld, 1996; Morris and Shin, 1998), for example, a central bank attempts to maintain a currency peg while speculators decide whether to attack the currency. A "regime change" occurs when enough speculators attack, forcing the bank to abandon the peg.

Similarly, in bank runs, a "regime change" happens when a large number of depositors lose confidence and withdraw their funds. If these withdrawals exceed the bank's liquid assets, the bank is forced to stop payments.

In models of self-fulfilling debt crises (Morris and Shin, 2004; Calvo, 1998; Cole and Kehoe, 1996, a "regime change" occurs when a borrower cannot secure refinancing due to a lack of creditor confidence. Atkeson (2000) uses a similar framework to describe riots, where the success of a rebellion against police forces depends on the size of the rebellion and the strength of the police.

Chapter 4

Comprehending Dynamic Global Games

In this chapter, we finally lay the foundations for the dynamic version of global games as discussed and deliberated by Angeletos, Hellwig, and Pavan (2007), specifically in the context of Regime Change. And we plan to do so by incorporating relevant assumptions to the erstwhile static framework understood in Chapter 3. Foremost assumption that we put in place is to allow the players to pursue actions repeatedly, across periods. In the view of the global games of Regime Change, this would essentially mean allowing the players to attack the status quo repeatedly. Hence, this assumption reinforces the realistic picture that there could be several attacks to the incumbent regime. Second of all, we let them accumulate private information across numerous periods, hence tracking the progression of a player's belief as the dynamic game goes on.

4.1 Setting up the Dynamic Framework

To this extent, consider the following setup from Angeletos, Hellwig, and Pavan (2007). Suppose a discrete time horizon, indexed by $t \in \{1, 2, 3, \cdots\}$ such that we hypothesize that the game would continue as long as the status quo is not abandoned, but ends as soon as the regime finally changes. Given a certain period t, let $R_t = 0$ reflect the event that the status quo is in place at the beginning of the period t, while $R_t = 1$ would then go on to describe the alternate event that it is abandoned at the start of the period t. Then, $a_{it} \in \{0, 1\}$ represents the action strategy for agent i, and by $A_t \in [0, 1]$, we describe the size of the

attack (total proportion of people choosing to invest) at period t. Fairly similar to the static framework where a successful attack was contingent upon the condition that the size of the attack must be sufficiently high, more formally that $A \geq \theta$, We continue to hold on to this condition in the purview of the dynamic game as well. We assume that if the status quo is in place at the beginning of a certain period t ($R_{t+1} = 1$), then the regime will only be abandoned if $A_t \geq \theta$. Likewise, the representation of the payoff functions in the dynamic setup draws inspiration from Equation 3.1 in the static model as a player $i's$ flow payoff for period t (conditional of $R_t = 0$) takes the form

$$\pi_{it} = a_{it}(R_{t+1} - c) \tag{4.1}$$

It then directly follows that his payoff from the entire game is given as the discounted sum of his payoffs across the individual periods, i.e $\Pi_i = \sum_{t=1}^{\infty} \rho^{t-1}(1 - R_t)\pi_{it}$, where $\rho \in (0, 1)$ is the discount factor. For a specific period, you may think about the discount factor as an arbitrary probability that the game continues to the next period.

Finally, the associated information structure to the dynamic setup is not much distant from the static case. Economic fundamentals, are summarized by θ and is drawn at the beginning of the game from the normal distribution $N(z, \frac{1}{\alpha})$, signifying the initial common prior belief of the players. Private information, however, is going to evolve across time periods. Specifically, across each period $t \geq 1$, every player i receives a private signal $\tilde{x_{it}} = \theta + \zeta_{it}$ about θ, where $\zeta_{it} \sim N(0, \frac{1}{\eta_t})$ is independently and identically distributed across i,and hence serially uncorrelated. Consequently, the player $i's$ history of private information acquisition up to period t is given by

$$\tilde{x_i}^t = \{\tilde{x}_{i\tau}\}_{\tau=1}^{t}$$

We also go by the hypothesis that the individual actions and the size of the past attacks are not observable, which is why the history of public information received up to period t would simply consist of information that the regime is still in place, whereas the private history for a player instead would be the sequence of private signals he has acquired in the past and the subsequent actions he has pursued in those periods.

Additionally, we define the sequence of private signals acquired up to period t by $\beta_t = \sum_{\tau=1}^{t} \eta_\tau$ and assume that $\infty > \beta_t \geq \frac{\alpha^2}{2\pi} \forall t$ and $\lim_{t\to\infty} \beta_t = \infty$. These assumptions encapsulate the two very fundamental features of a dynamic setting. First, that if we define a static game by letting the players act in a specific restricted period from the dynamic setting, then such a static game, must always possess a unique equilibrium for every t. Secondly, in the limit as $t \to \infty$, the private information becomes infinitely accurate . Now, while it is true that by hypothesizing player's payoffs per period to be completely independent on past or future actions, the model may fail to indulge interesting phenomenon like the effects of wealth accumulation in currency crises, we may be able to isolate information as a driver of coordination.

4.2 Reasonable conditions for Optimality

Before we can establish concrete conditions for stable equilibrium outcomes, we define certain terms that prove to be central to our analysis. Hence define $a_t : \mathbb{R}^t \to [0, 1]$ to be the action strategy profile for period t, while we hypothesize $a^t = \{a_\tau\}_{\tau=1}^{t}$ to be the strategy up to period t. In the same context, $a^\infty = \{a_\tau\}_{\tau=1}^{\infty}$, would reflect the complete strategy for the dynamic game. Having already assumed that private signals up to period t happen to be completely independent as well as identically distributed across the continuum of players, and obviously conditional on θ, for any strategy a^∞, the size of the attack only depends on the strength of the regime i.e θ. Finally let $p_t(\theta, a_t)$ be the probability that the status quo is abandoned in period t, when the state of the regime is θ, status quo is in place at the beginning of the period, and the governing action profile is a_t. Then, the cumulative distribution function of the posterior beliefs in period 1 can be parameterized by $\psi_1(\theta|\tilde{x_1})$; that for the subsequent periods $t \geq 2$ can be described through $\psi_t(\theta|\tilde{x}^t; a^{t-1})$ while continuing to hypothesize that the status quo is still in place ($R_t = 0$) and that agents have played in past periods according to a^{t-1}.

We must know that it is impossible for any player indulging in the competition to detect whether the play does or does not enforces equilibrium. And this is primarily because players cannot accurately predict their opponents' play in the periods leading up to a certain t. Furthermore, because payoffs to a player in a specific period is independent of what he himself did or what his opponent did in the past periods, the strategies will be sequentially rational if and only if they maximize the payoff to the player in that specific periods. Based on this hypothesis, a^{∞} defines an equilibrium if and only if the players do what is the best for themselves in each period. Henceforth, the following expressions must hold.

In the first period, $t = 1$, for private signal profile $\tilde{x_1}$, we have

$$a_1(\tilde{x_1}) \in \arg\max_{a\in[0,1]} \Big\{ \Big[\int p_1(\theta; a_1) d\psi_1(\theta|\tilde{x_1}) - c \Big] a \Big\}; \tag{4.2}$$

and similarly, for any subsequent period $t \geq 2$, for all $\tilde{x^t}$,

$$a_t(\tilde{x^t}) \in \arg\max_{a\in[0,1]} \Big\{ \Big[\int p_t(\theta; a_t) d\psi_1(\theta|\tilde{x^t}; a^{t-1}) - c]a \Big] \Big\} \tag{4.3}$$

It can be easily seen that Equations 4.2 and 4.3 define action profile which yield the highest payoff to the player in a certain period. Furthermore, the sequence of private information and that of private precisions, are defined recursively by x_t and β_t such that

$$x_t = \frac{\beta_{t-1}}{\beta_t} x_{t-1} + \frac{\eta_t}{\beta_t} \tilde{x_t} \tag{4.4}$$

and

$$\beta_t = \beta_{t-1} + \eta_t \tag{4.5}$$

with the private signal profile in the first period being $x_1 = \tilde{x_1}$, the precision profile being $\beta_1 = \eta_1$. Finally employing the standard Bayesian updating of the Gaussian prior beliefs we obtain that the strength of the regime θ assumes a normal distribution with mean $\frac{\beta_t}{\alpha+\beta_t} x_t + \frac{\alpha}{\alpha+\beta_t} z$ and precision $\beta_t + \alpha$. It follows that the x_t is a sufficient statistic for $\tilde{x}^t$ with respect to θ and hence, with respect to the event of regime change as well.

Note from the first period equilibrium condition given by Equation 4.2, it is clear that in any equilibrium, players play the first period of a dynamic in exactly the same way as the static

game discussed in Chapter 3. Therefore, from the Proposition 3.4.2, it is known that there exists a unique equilibrium in the first period where the players can only choose to attack at the first period, if at all they want to attack. Furthermore, this equilibrium is characterized in terms of thresholds for x_1 and θ. As for the subsequent periods, Consider the following Lemma by Angeletos, Hellwig, and Pavan (2007).

Lemma 4.2.1. *Any monotone equilibrium is characterized by a sequence* $\{x_t^*, \theta_t^*\}_{t=1}^{\infty}$ *with* $x_t^* \in \mathbb{R} \cup \{-\infty\}$, $\theta_t^* \in (0,1)$ *and* $\theta_t^* \geq \theta_{t-1}^*$ *for all* $t \geq 2$, *such that :*

i. at any period, an agent attacks if his acquired private signal $x_t < x_t^*$ *(certain threshold) and does not attack if* $x_t > x_t^*$

ii. the status quo is not abandoned in period t if and only if $\theta > \theta_{t-1}^*$

Proof. Angeletos, Hellwig, and Pavan (2007) make use of the Induction Principle to establish the proof. Employing the indifference result given by Equation 3.8 coupled with the Proposition 3.4.1 about the Uniqueness and Existence of Equilibrium outcomes, it is not very difficult to show that the Lemma is trivially validated for the first period. Following this observation, we make the inductive step by considering next any period $t > 2$ and hypothesizing that the Lemma also holds for any $\tau < t - 1$. Since players choose to go with a less dominating and rather safe action profile, (Not Attack, in the context of Regime Change) as private thresholds increase, we note that a_t does not increase with $\tilde{x^t}$. This observation coupled with the fact that the mass of attack $A_t(\theta)$ is non increasing in θ, only implies that either of the two following possibilities. Given the status quo is in place at the beginning of the current period t, either $A_t(\theta) < \theta$ and therefore $R_{t+1} = 0$ (Attack Unsuccessful), the game continues in the next period, for all $\theta > \theta_{t-1}^*$. Particular to this case, the state thresholds remain invariant across the considered periods i.e $\theta_t^* = \theta_{t-1}^*$. Or, there exists another situation, where we continue to have $\theta_t^* = \theta_{t-1}^*$ such that $A_t(\theta) < \theta$ (Attack still Unsuccessful) if and only if $\theta > \theta_t^*$ and the game terminates at the end of the

current period. In the former case, the posterior probability of regime change is 0 for all x_t and hence $x_t^* = -\infty$. In the latter case, the posterior probability of regime change is given by

$$\int p_t(\theta; a_t) d\psi_1(\theta|\tilde{x^t}; a^{t-1}) = \mathbb{P}(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = 1 - \frac{\Phi(\sqrt{\alpha+\beta_t}(\frac{\beta_t x_t + \alpha z}{\alpha+\beta_t} - \theta_t^*))}{\Phi(\sqrt{\alpha+\beta_t}(\frac{\beta_t x_t + \alpha z}{\alpha+\beta_t} - \theta_{t-1}^*))} \quad (4.6)$$

Because this posterior probability is continuous and strictly decreasing in relation to per period private signals x_t, it converges to 1 intuitively as beliefs about the regime strength are very low ($x_t \to -\infty$) and conversely, reduces to 0, as $x_t \to \infty$. Then from Intermediate Value Theorem, if the expression takes values at two extremes, it must take every value between these two extremes at some point. Hence, there exists an optimal $x_t^* \in \mathbb{R}$ such that:

- If $x_t = x_t^*$, then $\mathbb{P}(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = c$
- If $x_t < x_t^*$, $\mathbb{P}(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) > c$
- If $x_t > x_t^*$, $\mathbb{P}(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) < c$

In either case, the mass of attack, $A_t(\theta) < 1$ for all θ and hence $\theta_t^* < 1$ coupled with the fact that $\theta_t^* \geq \theta_1^* > 0$, implies that the $\theta_t^* \in (0, 1)$ for all t. This concludes the proof. □

Note that it is the strictly decreasing nature of the posterior probability which ensures that it is one to one and further c is achieved only at one point.

4.2.1 Implications of the Lemma

Given the underlying setup, there is no doubt that if the status quo is in place in a certain period, then it must also be in place in the preceding periods. Otherwise, the dynamic game, would have simply ended earlier. More formally, In any period, $t \geq 2$, if the status quo is in place, then it makes perfect sense to think that the actual realization of the status quo must have been stronger than the threshold of the previous period, i.e $\theta > \theta_{t-1}^*$, otherwise players

would not only have had the incentive to attack but it would also have been successful leading to the end of the game. Therefore, the sequence $\{\theta_t^*\}$ is non decreasing. In contrast, the sequence for private information thresholds $\{x_t^*\}$ does not follow a strict pattern and is seen as non monotonic in general. Imperatively, There might some periods when there is negative perception about the strength of the regime, $x_t^* > -\infty$ and only some players choose to attack while there might also be some periods where $x_t^* = -\infty$ and no player attacks.
Having arrived at the conclusion, that at equilibrium, players do play the first period of the dynamic game in exactly the same fashion as the static game, we acknowledge that such cannot be said for the subsequent periods. The fact that the sequence $\{\theta_t^*\}$ is non decreasing enables us to write $\theta_{t-1}^* \geq \theta_1^* > 0$. This further gives $\theta > \theta_1^*$, meaning that there would always be equilibria in which nobody attacks in a period $t \geq 2$. In particular, there always exists an equilibrium, in which an attack would only happen in the first period but never thereafter. When such an outcome is the unique equilibrium outcome, the potential to take repeated actions against the regime would not be any more interesting or unexpected than the anticipatory outcome.

4.2.2 Equilibrium with Multiple attacks

Having shown what really transpires when the case of an attack only being plausible in the first period happens to be unique robust outcome, we find it increasingly useful to assess the conditions under which there may also be other equilibria where an attack is not just restricted to the first period. And this is precisely, where we resort to the study put forward by Angeletos, Hellwig, and Pavan (2007).
From Lemma 4.2.1, we refrain from considering $x_t^* = \infty$, or the case where every player chooses to attack. Reason being the fact that irrespective of how accurate the private information goes to become, the status quo will always survive for $\theta > 1$ and hence it will

be dominant not to attack. Therefore, We are only interested in equilibria in which $x_t^* \in \mathbb{R}$. The size of the attack is then given by

$$A_t(\theta) = Pr(x_t \leq x_t^* | \theta) = \Phi(\sqrt{\beta_t}(x_t^* - \theta_t)) \tag{4.7}$$

which is continuous and strictly decreasing in θ, while the posterior probability of regime change for a player with private history statistic x_t is given by the expression 4.6, which is continuous and strictly decreasing in x_t, with the regular assumption of $\theta_t^* > \theta_{t-1}^*$. Finally, it follows that if there is any equilibrium, in which an attack happens in a period t, θ_t^* and x_t^* solve

$$\theta_t^* = A_t(\theta_t^*)$$

and

$$\mathbb{P}(\theta \leq \theta_t^* | x_t^*, \theta > \theta_{t-1}^*) = c$$

which is equivalent to writing,

$$\theta_t^* = \phi(\sqrt{\beta_t}(x_t^* - \theta_t^*)) \tag{4.8}$$

and

$$1 - \frac{\Phi(\sqrt{\alpha + \beta_t}(\frac{\beta_t}{\beta_t + \alpha} x_t^* + \frac{\alpha}{\beta_t + \alpha} z - \theta_t^*))}{\Phi(\sqrt{\alpha + \beta_t}(\frac{\beta_t}{\beta_t + \alpha} x_t^* + \frac{\alpha}{\beta_t + \alpha} z - \theta_{t-1}^*))} = c \tag{4.9}$$

These conditions 4.8 and 4.9 are the dynamic game analogues of the conditions 3.7 and 3.9 which we discussed in the static framework in Chapter 3. While the condition 4.8 only implies that the equilibrium size of the attack is nothing but the critical size which would cause regime change given the fundamentals are θ_t^*, the condition 4.9 contrarily ensures that the agent must be indifferent between attacking and not attacking if and only if he receives the private information derived from the signal x_t^*.

4.3 Alternate Expressions of Equilibrium

Following analysis inspired by Angeletos, Hellwig, and Pavan (2007). For convenience, the equilibrium conditions defined in equations 4.8 and 4.9 can also be written down in terms of

several functions described as follows:

Consider the functions $u : \mathbb{R} \times [0,1] \times \bar{\mathbb{R}} \times \mathbb{R}^2_+ \times \mathbb{R} \to [-c, 1-c]$, $X : [0,1] \times \mathbb{R}_+ \to \bar{\mathbb{R}}$, and $U : [0,1] \times \bar{\mathbb{R}} \times \mathbb{R}^2_+ \times \mathbb{R} \to [-c, 1-c]$ which are defined as follows :

$$u(x, \theta^*, \theta^*_{-1}, \beta, \alpha, z) \equiv \begin{cases} 1 - \frac{\Phi(\sqrt{\alpha+\beta}(\frac{\beta}{\beta+\alpha}x + \frac{\alpha}{\beta+\alpha}z - \theta^*))}{\Phi(\sqrt{\alpha+\beta}(\frac{\beta}{\beta+\alpha}x + \frac{\alpha}{\beta+\alpha}z - \theta^*_{-1}))} - c & if \theta^* > \theta^*_{-1} \\ -c & if \theta^* \leq \theta^*_{-1} \end{cases} \tag{4.10}$$

The function, $u(x, \theta^*, \theta^*_{-1}, \beta, \alpha, z)$ is the net payoff from attacking to a player who observes a private information statistic x_t when it is known that the status quo is not abandoned at the beginning of that specific period. Now, recall the first equilibrium condition we described in Equation 4.8. Solving this expression for x_t gives us the following alternate representation,

$$X(\theta^*, \beta) \equiv \theta^* + \frac{1}{\sqrt{\beta}}\Phi^{-1}(\theta^*) \tag{4.11}$$

which works on the hypothesis that if the players attack in the period t if and only if their acquired private signal in that period, $x_t \leq x^*$, then the attack is only going to be successful in that period, $A_t(\theta) \geq \theta$ if and only if this actual strength is less than the state threshold, i.e $\theta \leq \theta^*$. Lastly, we also characterize the function :

$$U(\theta^*, \theta^*_{-1}, \beta, \alpha, z) \equiv \begin{cases} \lim_{x \to -\infty} u(x, \theta^*, \theta^*_{-1}, \beta, \alpha, z) & if \theta^* = 0 \\ U(X(\theta^*, \beta), \theta^*, \theta^*_{-1}, \beta, \alpha, z) & if \theta^* \in (0,1) \\ \lim_{x \to \infty} u(x, \theta^*, \theta^*_{-1}, \beta, \alpha, z) & if \theta^* = 1 \end{cases} \tag{4.12}$$

as the net payoff from attacking to an average player who observes the private signal threshold defined previously be $X(\cdot)$, obviously conditional on the assumption that the status quo is in place at the beginning of the period, i.e $\theta > \theta^*_{-1}$.

Now, substituting the alternate expression for the threshold, x^*_t into the equilibrium conditions defined by Equation 4.9, we finally obtain :

$$U(\theta^*_t, \theta^*_{t-1}, \beta_t, \alpha, z) = 0 \tag{4.13}$$

Condition 4.13 would continue to be a very crucial characterization result for most proofs and outcomes we discuss hereafter. To be more precise, this condition is essentially nothing

but the indifference criterion for choosing between attacking and not attacking for a marginal player who is observing a private statistic x_t in a period t for $t \geq 2$. We do not include the first period $t = 1$ here, since in the first period, the status quo has never been attacked in the past. Consequently, the indifference criterion for specifically the first period comes to be $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$. Note that $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$ is exactly identical to the payoff to any marginal player in the static framework.

In conclusion, we are finally in the position to prove an alternate and more accessible condition for equilibrium, making use of some very important functions we have obtained just previously.

4.4 Embracing Monotonicity

Proposition 4.4.1 (Angeletos, Hellwig, and Pavan, 2007)**.** *The strategy* $\{a_t(\cdot)\}_{t=1}^{\infty}$ *is a monotone equilibrium if and only if there exists a sequence of thresholds* $\{x_t^*, \theta_t^*\}_{t=1}^{\infty}$ *such that*

i. for all periods, $a_t(\tilde{x}^t) = 1$ *if* $x_t < x_t^*$ *and* $a_t(\tilde{x}^t) = 0$ *if* $x_t > x_t^*$

ii. for $t = 1$, θ_1^* *solves* $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$ *and* $x_1^* = X(\theta_1^*, \beta_1)$*;*

iii. for any period except the first one, either $\theta_t^* = \theta_{t-1}^* > 0$ *and* $x_t^* = -\infty$*, or* $\theta_t^* > \theta_{t-1}^*$ *is a solution to* $U(\theta_t^*, \theta_{t-1}^*, \beta_t, \alpha, z) = 0$ *and* $x_t^* = X(\theta_t^*, \beta_t)$*.*

A monotone equilibrium always exists.

Proposition 4.4.1 renders us a recursive algorithm which can be used to obtain an entire set of monotone equilibria for a dynamic game close to the nature we have been discussing so far. We start with the first period, $t = 1$, and denote θ_1^* as the unique solution to $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$. The expression would always admit the solution for the first period since it has already been established that the first period of our model is exactly identical to the static framework, in that it will always admits a monotone equilibrium.

We subsequently proceed to the next period, $t = 2$. If the utility equation $U(\theta_2^*, \theta_1^*, \beta_2, \alpha, z) = 0$ admits a solution, then let θ_2^* be such a solution. If does not have any solution whatsoever, then we may simply set $\theta_2^* = \theta_1^*$. And we continue in the same fashion for all the subsequent periods.

This particular set of sequences $\{\theta_t^*\}_{t=1}^{\infty}$ obtained by the means of recursion, coupled with the associated sequence of private signal threshold $\{x_t^*\}_{t=1}^{\infty}$ gives the set of monotone equilibria. Furthermore, we note that the algorithm works on the principle that the equilibrium learning can be depicted in a simple form of a truncation in the support of beliefs about θ.

To conclude, there always exist at least one monotone equilibria which comes as a solution to the first period payoff function but the question of whether the same can be said about other periods is interesting to study. Hence, let us assess the characteristics of the U function by the means of the following lemma and try to identify the conditions under which there are other solutions to the Indifference criterion 4.13, mentioned earlier.

Lemma 4.4.1 (Angeletos, Hellwig, and Pavan, 2007). *i. The function $U(\theta^*, \theta_{-1}^*, \beta, \alpha, z)$ is continuous in all its arguments, non monotonic in θ^* when $\theta_{-1}^* \in (0, 1)$, and strictly decreasing in θ_{-1} and z for $\theta_{-1}^* < \theta^*$. Furthermore, for all $\theta_{-1}^* < 1$ and $\theta^* > \theta_{-1}^*$, $\lim_{\beta\to\infty} U(\theta^*, \theta_{-1}^*, \beta, \alpha, z) = \theta_\infty - \theta^*$, where $\theta_\infty \equiv 1 - c$.*

ii. Let $\hat{\theta}_t$ be the unique solution to the $U(\hat{\theta}_t, -\infty, \beta_t, \alpha, z) = 0$. A solution to the indifference criterion 4.13 would only exist if $\theta_{t-1}^* < \hat{\theta}_t$ and must be necessarily bounded above by $\hat{\theta}_t$.

iii. If $\theta_{t-1}^* > \theta_\infty$, a solution to the indifference criterion 4.13 does not exist even for sufficiently precise sequence of private signals, i.e β_t.

iv. If $\theta_{t-1}^* < \theta_\infty$, a solution to the indifference criterion 4.13 would exist even for sufficiently precise sequence of private signals (β_t).

v. If θ_{t-1}^* is the highest solution to the criterion 4.13 for period $t-1$, there exists $\underline{\beta} > \beta_{t-1}$ such that 4.13 admits no solution for any period $\tau \geq t$ such that $\beta_\tau < \underline{\beta}$.

4.4.1 Interpretation of Lemma

The function U is clearly non monotonic in θ^* given the hypothesis that $\theta^*_{-1} \in (0,1)$. As the state threshold θ^* increases, the private signal threshold x^* also increases, according to the expression $x^* \equiv \theta^* + \frac{1}{\sqrt{\beta}}\Phi^{-1}(\theta^*)$. Players will attack in period t if and only if they acquire $x_t \leq x^*$, leading to the status quo being abandoned if $A_t(\theta) \geq \theta$, i.e, $\theta \leq \theta^*$.

Following this logic, When the period t state threshold is less than the previous period threshold, $\theta^* < \theta^*_{-1}$, it results in a low private signal threshold x^*. Consequently, players become less aggressive, resulting in smaller attack sizes than θ for all $\theta > \theta^*_{-1}$. Because the perceived probability of regime change for each player is low, we have : $U(\theta^*, \theta^*_{-1}, \cdot) = -c$ for $\theta^* < \theta^*_{-1}$.

Conversely, when the current state threshold is higher than the one in period $t-1$, $\theta^* > \theta^*_{-1}$, the private signal threshold x^* is so high that there is a significant possibility of regime change, assuming the status quo is maintained at the beginning of the period. This implies, $U(\theta^*, \theta^*_{-1}, \cdot) > -c$ for $\theta^* > \theta^*_{-1}$.

Additionally, as $\theta^* \to 1$, from the expression, $x^* \to \infty$, and therefore, the probability that any marginal player assigns to the event that $\theta > 1$ converges to 1. This means that the probability they assign to status quo being abandoned is 0. Thus, $U(\theta^*, \theta^*_{-1}, \cdot) \to -c$ as $\theta^* \to 1$.

To diagrammatically represent this non monotonic behaviour exhibited by the payoff function U with respect to θ^*, we consider the plot 4.1. In this plot, the value is $-c$ for $\theta^* < \theta^*_{-1}$, it subsequently increases as θ^* increases, and eventually starts decreasing and, finally converges to -c as $\theta^* \to 1$. Any intersection of this plot with the horizontal axis (reflecting θ^*) signifies a solution to the indifference criterion defined through expression 4.13. The dashed line plots the payoff to a marginal player in the static framework, where players can only attack in the period t. This function is monotonic with respect to θ^*, when the players have a sufficiently precise private signal β_t in period t. This monotonic behaviour of the payoff function U in the static framework ensures that we have a unique equilibrium. In the dynamic game, the

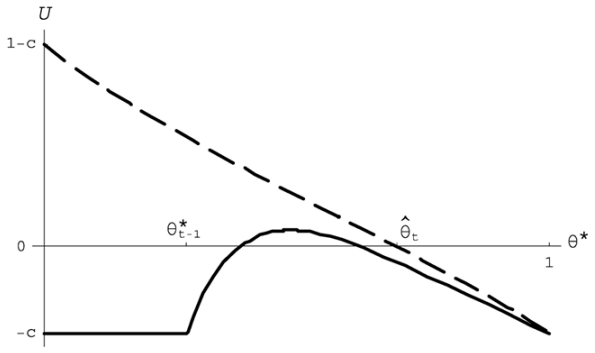

Figure 4.1: Non Monotonicity in θ^*(Angeletos et al.(2007)

absence of monotonicity leaves open the possibility for multiple equilibria.

Examining $U's$ response to variations in prior mean z, we find that U negatively relates to z. As players' beliefs in the strength of the status quo increases, their incentive to attack decreases, leading to a lower probability of regime change for any monotone strategy, and lower payoff gain to a marginal player from choosing to attack.

Regarding relationship of U with respect to θ^*_{-1}, as θ^*_{-1} increases, it implies that the players were aggressive in the periods prior. If the status quo still stands at the beginning of the current prevailing period, players' posterior beliefs about the regime's strength increase, reducing the incentive to attack and thus lowering U. Therefore, U also decreases with θ^*_{-1}. Using these bi variate relationships, it is evident that the payoff from attacking for a marginal player in any period except the first one, is always lower than the payoff in the first period, i.e $U(\theta^*, -\infty, \beta_t, \alpha, z)$ given the private signal precision is β_t. This is also evident from the plot.

To better comprehend iii. and iv., note that as the private signals become perfectly accurate, the impact on players' posterior beliefs that the status quo is maintained, diminishes for any $\theta_t > \theta^*_{t-1}$ and by consequence, the payoff difference for a marginal player attacking in the first period versus any subsequent period also diminishes.

This coupled with the observation that, $U(\theta^*, -\infty, \beta_t, \alpha, z) \to \theta_\infty - \theta^*$ as $\beta_t \to \infty$, only ensures that for sufficiently high β_t, the indifference criterion would admit at least one

solution if $\theta^*_{t-1} < \theta_\infty$ and no solution if $\theta^*_{t-1} > \theta_\infty$, where $\theta_\infty = \lim_{t\to\infty} \hat{\theta}_t$ is the limit of the equilibrium threshold in the static game for $\beta \to \infty$

To understand v. consider the scenario where the largest possible attack in one period is unsuccessful. This increases the posterior beliefs about the regime's strength, deterring further attacks unless significant new private information is received. This feature highlights the importance of new private information as a necessary criterion for further attacks to become possible after period 1. Whether this is a sufficient condition depends on the prior mean z.

For a low (aggressive) prior, players start with beliefs less supportive of the regime. As the game progresses and the regime survives the initial attack, players would tend to revise their beliefs upwards, decreasing aggression in subsequent periods. Specifically for low z, the equilibrium threshold $\hat{\theta}_t$ decreases over time, deterring further attacks and leading to a unique monotone equilibrium where all agents follow a strategy which leads with no attacks after the first period.

Conversely, for a high z (lenient) prior, players start with strong support for the status quo, leading to a higher initial threshold for attacking. Over time, $\hat{\theta}_t$ increases, counteracting the positive belief induced by the regime's survival, creating conditions for possible new attacks and sustaining multiple equilibria, including the scenarios where attacks occur in multiple periods.

4.4.2 An Instance of Multiplicity

Following example from Angeletos, Hellwig, and Pavan (2007). Consider a scenario when $\theta^*_1 < \theta_\infty$, (i.e the prior mean z is sufficiently high by the logic we have established in the last section), then the second attack necessarily becomes very much possible conditional on the private signal β_t being large enough. Take the following U vs θ^* plot given by Figure 4.2 to illustrate multiplicity in equilibrium outcomes across such instances.

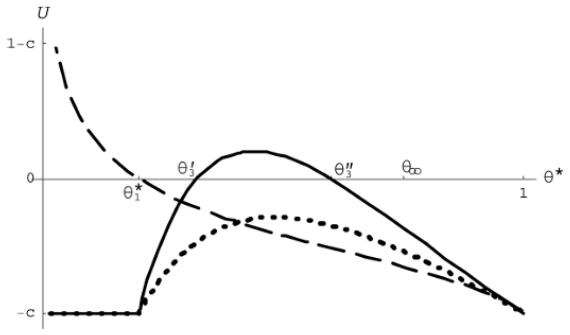

Figure 4.2: Angeltos et al. (2007)

The dashed line plots the payoff of the marginal player in the first period, $t = 1$ as its intersection with the horizontal axis defines $\theta_1^* < \theta_\infty$. Her payoff in the period $t = 2$ and $t = 3$ is represented by the dotted and the solid line respectively. As evident from the diagram, the private information acquisition in the second period is low enough that no player attacks in the period 2 and the status quo prevails. In contrast, The acquisition of private information across the third period becomes high enough that a new attack is possible. Therefore, for this particular example, there exists at least three equilibria : the first one, in which the state threshold remains the same across all three periods, i.e $\theta_t^* = \theta_1^*$ for all t, another in which $\theta_2^* = \theta_1^*$ and $\theta_t^* = \theta_3'$ for all $t \geq 3$, and a third in which $\theta_2^* = \theta_1^*$ and $\theta_t^* = \theta_t^* = \theta_3''$ for all $t \geq 3$, where θ_3' and θ_3'' correspond to the two intersections of the solid payoff curve with the horizontal axis.

In conclusion, we note that if z is sufficiently high, any solution to indifference criterion, is strictly less than θ_∞ in all periods, which ensures that a new attack eventually becomes possible after an unsuccessful one.

Chapter 5

Rationally Inattentive Acquisition of Information

5.1 A Non Technical Introduction to Rational Inattention

In real-world scenarios, decision-makers often face overwhelming amounts of data, making it impractical and costly to process all available information. As a result, they must focus selectively on certain details while ignoring others. This selective attention underpins rational inattention theory, where agents allocate their limited attention to the most relevant information to maximize expected utility.

Introduced by Sims (2003), Rational Inattention has significantly impacted economics and finance. In macroeconomics, it explains why consumers might not immediately react to monetary policy changes due to the costs of staying informed. It also clarifies patterns in asset price movements, as investors selectively process news and financial reports, and sheds light on investor overreaction to public news. In Regime Change Models, rational inattention helps explain how agents decide how much information to gather about potential events like currency peg abandonments or bank runs.

5.2 Foundations - The Rational Inattention Game

Recall the parametric class of two player two action games $G(\theta)$ that we had worked on in Section 2.2, where the payoffs at stake factor on the realizations of the economic fundamental, θ. For the sake of convenience, we let the two players be parameterized by i and j, so that we refer to the player i as the generic player and let $j \neq i$ be his opponent. For a generic player i, let A^i be the action- set comprising of two possible action strategies - "Invest" and "Not Invest". Then the set of all action profiles would then be given by $A = A^1 \times A^2$. Consequently, the set of all possible economic fundamentals be given as Θ of which θ is a generic element. We obviously continue with the grounded assumption that θ is not common knowledge, and moreover, there is a full support common prior $\Omega \in \Delta(\Theta)$. Additionally, let the payoffs to i be given by $u_i : A \times \Theta \to \mathbb{R}$, so that $u_i = (a_i, a_j, \theta)$ would go on to represent the utility to player i for choosing a_i, while his opponent chooses a_j and the prevailing fundamental is θ. Since players happen to be uncertain about the kind of payoff structure they are dealing with, they have an incentive to acquire more knowledge about the parameter pair (a_j, θ). In more precise terms, they only have access to a noisy private information surrounding the realization which eventually determines their optimal strategy, or the conditional action choice probability $p_i(\cdot|a_j, \theta) \in \Delta(A^i)$.

5.2.1 Curious Case of Flexible Information Structure

In the study of coordination games and rational inattention, the nature of the information structure plays a crucial role in determining players' strategic behavior and equilibrium outcomes. Global games, typically employ an additive information structure where each player's private signal is modeled as a sum of the true state and some noise. This additive noise framework facilitates the formation of higher-order beliefs, where players sequentially update their

beliefs based on the noisy signals received from the environment, ultimately leading to a unique equilibrium (under special conditions). Conversely, rational inattention (RI) theory introduces a more flexible information structure, allowing agents to optimally choose the amount and type of information they acquire, constrained by their information-processing capacity. This flexibility is captured through entropy-based costs and mutual information constraints, enabling agents to focus their limited attention on the most relevant aspects of the state. Unlike the fixed and exogenous noise in global games, the endogenous choice of information in RI results in a dynamic and often more complex equilibrium formation, as agents continuously adjust their information acquisition strategies to optimize decision-making under uncertainty.

5.2.2 Conditional Choice probabilities as a response to Information Signals

Having acknowledged the features of the RI Framework, there are certain expectations from a model that incorporates these assumptions from Behavioral Economics. Therefore we hypothesize an arbitrary player to pursue a certain action as if he is choosing a likelihood of opting for that action conditional on a_j and θ, $p_i(\cdot|a_j, \theta)$, for each (a_j, θ). Hoshino (2018) refers to these conditionals as 'strategies' in the game and denotes them through $\mathcal{P}_i = (p_i(\cdot|a_j, \theta))_{a_j,\theta}$. We would follow suit. This strategy is reflective of what information to acquire, and also takes into account the mistakes involved in computation. While one has no influence about the actions pursued by one's opponents(at least we assume this to be the case), one might still acquire the marginal conditional, given by $p_{j\theta} \in \Delta(A^j \times \Theta)$ through exogenous sources. It is useful to note here that the marginal (joint distribution) on a certain state space Θ must coincide with the player's full support common prior given by Ω. Mathematically, it should be the case that $\text{marg}_\Theta p_{j\theta} = \Omega$.

Finally, player $i's$ strategy $\mathcal{P}_i$ as well as his marginal $p_{j\theta}$ together can be employed to define

a joint $p \in \Delta(A \times \Theta)$ so that the player's expected payoff becomes :

$$\mathbb{E}_p\big[u_i(a_i, a_j, \theta)\big] = \sum_{a_j,\theta} p_{j\theta}(a_j, \theta) \sum_{a_i} p_i(a_i|a_j, \theta) u_i(a_i, a_j, \theta) \tag{5.1}$$

Try to interpret this payoff as product of two sums. While the first sum corresponds to the sum of the marginal probabilities as a_j and θ varies, the second sum is that of conditional probabilities of choosing an action (strategies for a player i), weighted by the utility to the same player contingent upon the chosen strategies.

5.3 Cognitive Costs of Information

Rational inattention models use a cost function to quantify the cognitive effort needed to process information. This function, often based on information theory measures like the Kullback-Leibler divergence, captures the trade-off between information precision and cognitive resources. To keep these ideas intact, let us consider the following example by Hoshino (2018), where the cost of information acquisition for a player i is represented as $c_i : \Delta(A \times \Theta) \to \mathbb{R}_+$ furthering the cost $c_i(p) \geq 0$ of the joint choice probability p. Essentially, contingent upon his marginal strategy being $p_{j\theta} \in \Delta(A^j \times \Theta)$, player i would incur the cost $c_i(p_i)$ to choose a strategy $\mathcal{P}_i$ that induces p_i. Without the loss of generality, $c_i(p_i)$, can be alternatively written as $c_i(\mathcal{P}_i, p_{j\theta})$. Now, we want our cost functional to satisfy certain central properties which rule out exclusive and vague settings, and thereby provide normative support for using a general non parametric form rather than a very restrictive one. These properties are as follows :

Definition 5.3.1 (Monotonicity of Cost)**.** *(Hoshino, 2018) Given a player i's marginal strategy $p_{j\theta}$, i's information cost functional c_i is said to be monotone if for any two arbitrary strategies $\mathcal{P}_i, \mathcal{Q}_i$, acquiring more information about $\mathcal{P}_i$ is more beneficial than knowing more*

about $\mathcal{Q}_i$*:*

$$\sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} p_i(a_i|a_j,\theta) u_i(a_i,a_j,\theta) > \sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} q_i(a_i|a_j,\theta) u_i(a_i,a_j,\theta) \tag{5.2}$$

or simply, $c_i(p_i) > c_i(q_i)$

This property also reinforces the Consistency principle (Blackwell, 1953) which states that more informative signals should not cost less. The second fundamental property only strengthens the argument by saying that there is no cost to uninformative conditionals :

Definition 5.3.2 (Uninformative Information Strategy)**.** *Player i′s strategy* $\mathcal{P}_i$ *is an uninformative information strategy if given one's marginal* $p_{j\theta}$*, the cost of acquiring* $\mathcal{P}_i$ *is 0, i.e* $c_i(\mathcal{P}_i, p_{j\theta}) = 0$

Undoubtedly, if we hope to eventually feed a similar cost functional to a dynamic setup, we must incorporate a set of numerous additional properties. We will describe these later.

5.4 Objectives of the RI Game

As soon as we incorporate costs of acquiring information to a baseline global game framework, we are essentially dealing with a Rational Inattention Game - say RI game for convenience. Such an RI Game is denoted by $\langle G, c \rangle$. As per our intuition, the optimization problem faced by rationally inattentive agents in this RI game $\langle G, c \rangle$ involves choosing an information acquisition strategy $\mathcal{P}_i$ that balances the expected benefits of better decision-making against the cognitive costs of acquiring and processing information. Therefore, putting it all together, a player's objective can then be represented as: Given a marginal $p_{j\theta}$,

$$\max_{\mathcal{P}_i} \mathbb{E}_p\big[u_i(a_i,a_j,\theta)\big] - c_i(\mathcal{P}_i, p_{j\theta})$$

$$\implies \max_{\mathcal{P}_i} \sum_{a_j,\theta} p_{j\theta}(a_j,\theta) \sum_{a_i} p_i(a_i|a_j,\theta) u_i(a_i,a_j,\theta) - c_i(\mathcal{P}_i, p_{j\theta}) \tag{5.3}$$

5.5 Amending Bayesian Correlated Equilibrium

With the aforementioned objective in place for the RI game $\langle G, c \rangle$, we recognize the importance of each player's strategies being informed and profit maximizing conditional on the joint probabilities over the opponents' behavior as well as the state. The solution concept which is usually brought in place is not much distinct from the Bayesian Correlated Equilibrium introduced by Bergemann and Morris (2016) except for the fact that we also do have to bring information costs into consideration.

Definition 5.5.1 (Bayesian Correlated Equilibrium with Information Costs (Hoshino, 2018))**.** *An Equilibrium of the RI Game $\langle G, c \rangle$ is a joint conditional $p^* \in \Delta(A \times \Theta)$ such that :*

i. Marginal conditionals are consistent with the full support initial prior Ω, i.e $marg_\Theta p^* = \Omega$

ii. For each player i, their strategy $\mathcal{P}_i^* = (p^*(\cdot|a_j,\theta))_{a_j,\theta}$ is optimal to the marginal $p_{j\theta}^* = marg_{A^j \times \Theta} p^*$. Hence for any strategy $\mathcal{P}_i = (p_i(\cdot|a_j,\theta))_{a_j,\theta}$, we must have:

$$\begin{aligned} &\sum_{a_j,\theta} p_{j\theta}^*(a_j,\theta) \sum_{a_i} p^*(a_i|a_j,\theta) u_i(a_i,a_j,\theta) - c_i(\mathcal{P}_i^*, p_{j\theta}^*) \\ &\geq \sum_{a_j,\theta} p_{j\theta}^*(a_j,\theta) \sum_{a_i} p_i(a_i|a_j,\theta) u_i(a_i,a_j,\theta) - c_i(\mathcal{P}_i, p_{j\theta}^*) \end{aligned}$$

5.6 The Case of Shannon's Entropy Reduction Functional

In the context of Rational Inattention, the information acquisition costs are often modeled using Shannon's entropy reduction cost functional. This approach is grounded in information

theory, where entropy measures the uncertainty or unpredictability associated with a random variable. Mathematically, this cost functional takes the form : (Hoshino, 2018)

$$c_H(p) := \sum_{a_j,\theta} p(a_j,\theta) \sum_{a_i} \big(\log p_i(a_i|a_j,\theta) - \log p_i(a_i)\big) \tag{5.4}$$

Here, $p_i(a_i) = p_i(a_i|a_j,\theta)$ represents the unconditional probability of choosing a certain action a_i. Let the term following the second summation in the above expression be denoted as H. This term H is essentially a measure of how far the conditional probability of choosing a certain action $p_i(\cdot|a_j,\theta)$ is from the unconditional probability of doing the same $p_i(\cdot)$ using the difference in terms of the logarithms. When it is the case that the information costs $c = c_H$, player $i's$ optimal strategies (or conditionals) takes the form of logit expression given by Matějka and McKay (2015) as: For reciprocal $\beta = \frac{1}{\lambda}$, λ being the scale of information acquisition cost.

$$p_i(a_i|a_j,\theta) = \frac{p_i(a_i)\exp\beta u_i(a_i,a_j,\theta)}{p_i(1)\exp\beta u_i(1,a_j,\theta) + p_i(0)\exp\beta u_i(0,a_j,\theta)} \tag{5.5}$$

These strategies defined by the logit expression do factor on the unconditional probabilities $p_i(\cdot)$, but they are endogenously determined at equilibrium.

While most theoretical RI games model information acquisition through Shannon's Entropy Reduction, this approach can yield inconsistent or even implausible results, as noted by Dean and Neligh (2023) and Fosgerau and Jiang (2019). Hence, following Hoshino (2018), we opt for a non-parametric form for the cost functionals instead of limiting ourselves to Entropy Reduction. This choice ensures flexibility while adhering to essential properties. Moreover, there's no compelling reason to constrain factors like cognitive costs or market research solely to Shannon's framework.

5.7 Plausible Extensions in the Dynamic Framework

In the dynamic global game model introduced by Angeletos, Hellwig, and Pavan (2007) and discussed in detail in Section 4.1, players are allowed chances to attack a regime across

multiple periods. The status quo regime has an unknown strength parameterized by θ, and players acquire signals about θ over time. Broadly, their strategies factor on these multi-temporal signals as well as on the observed history of attacks. In this section, we attempt to introduce information costs to this setup to further enhance our understanding of how players gather and utilize information over multiple periods while incurring cognitive costs. To this extent, we do not restrict ourselves to the use of Shannon's Entropy Reduction but rather go by the non parametric form proposed in Hoshino (2018).

5.7.1 Setup

We build our standard hypothesis on the similar lines as Angeletos, Hellwig, and Pavan (2007). There are N players indexed by i and each player has access to an identical strategy set,i.e $a_{it} \in \{0, 1\}$. $A_t \in [0, 1]$ describes the size of the attack in a period t, which is only successful iff $A \geq \theta$, given that the status quo is maintained at the beginning of the period t,$R_t = 0$. The regime's strength is represented by a fundamental θ, which is unknown to the players and drawn from $N(z, \frac{1}{\alpha})$. Across each period $t \geq 1$, a player i receives a private signal $\tilde{x_{it}} = \theta + \zeta_{it}$ about θ, where $\zeta_{it} \sim N(0, \sigma_t^2)$, precision being ρ_t, only at a certain cost. We define this information cost, $c(\rho_t)$, to be non parametric, factoring on the precision ρ_t of the private signal x_{it} acquired by any player i in a period t. We further assume that $c(\rho_t)$ is increasing in ρ_t, i.e, more precise (informative) signals cost more (Blackwell, 1953). Specifically, we take $c(\rho_t)$ to be a strictly convex function given by

$$c(\rho_t) = \varpi(\rho_t - \rho_0)^2 \tag{5.6}$$

where $\varpi > 0$ is a parameter that penalizes deviations from a baseline precision ρ_0 and in addition, we also have

$$\frac{\partial^2 c(\rho_t)}{\partial \rho_t^2} > 0$$

5.7.2 Dynamic Strategies

At each period, a player i must decide the level of signal precision ρ_t they want to acquire at each period t, and simultaneously also whether they should attack based on their posterior beliefs and the signals they receive. Furthermore, the decision at each period depends on the expected future payoffs, and the history of signals received up to a certain period t, given by $\tilde{x}^t$. To this context, let $\Upsilon_t(\theta, \tilde{x}^t)$ represent the value function for a player at a period t, θ being the true regime strength. In the original setup, this value function is updated using Bayesian Inference. However, with the introduction of non-parametric costs, Υ needs to incorporate the cost of acquiring information. Therefore, we must have :

$$\Upsilon_t(\theta, \tilde{x}^t) = \max_{\rho_t}\{\mathbb{E}[U_t(x_{it}, \theta_t, \tilde{x}^t)|\rho_t] - c(\rho_t) + \delta\mathbb{E}[\Upsilon_{t+1}(\theta, \tilde{x}^{t+1})|\tilde{x}^t, \rho_t]\} \tag{5.7}$$

where $U_t(x_{it}, \theta_t, \tilde{x}^t)$ as inspired from Angeletos, Hellwig, and Pavan (2007), is the immediate payoff from attacking or not attacking in period t. Similarly, δ or the discount factor, governs the probability of moving to the next period, and is taken product with the expected value in the subsequent periods. Also, given that the signal is normally distributed, the expected utility can often be approximated as increasing in signal precision, but rate of increase diminishes because the marginal benefit of reducing uncertainty becomes small.

Given the recursive nature of the game as well as intuition to some extent, players would update their strategies based on their trade-off between the cost of acquiring more precise information and the expected future payoffs. Therefore, the optimal choice of signal precision ρ_t^* should maximize the value function Υ_t in each period t, taking into account both the immediate and future costs and benefits. Mathematically, we should imply that the optimal strategy across periods is derived from the function

$$\Upsilon_t(\theta, \tilde{x}^t) = \max_{\rho_t}\{\mathbb{E}[U_t(x_{it}, \theta_t, \tilde{x}^t)|\rho_t] - c(\rho_t) + \delta\mathbb{E}[\Upsilon_{t+1}(\theta, \tilde{x}^{t+1})|\tilde{x}^t, \rho_t]\} \tag{5.8}$$

And the optimal signal precision at period t is determined by :

$$\rho_t^* = \arg\max_{\rho_t}\{\mathbb{E}[U_t(x_{it}, \theta_t, \tilde{x}^t)|\rho_t] - c(\rho_t) + \delta\mathbb{E}[\Upsilon_{t+1}(\theta, \tilde{x}^{t+1})|\tilde{x}^t, \rho_t]\} \tag{5.9}$$

5.8 Reaching Equilibrium Uniqueness

In this section, we finally obtain conditions for uniqueness in equilibrium outcomes. We have had a sound discussion in the Introduction concerning the significance of such robustness in driving comparative statics as well as policy implications. However, before we draw any inferences about the description of equilibrium outcomes in this novel setting, we first claim that if we take the marginal benefit of increasing precision to be identical for all players in the setup, then the convex nature of information costs would drive all players to choose the same level of signal precision in equilibrium. Hence consider the optimality condition for signal precisions, given by Equation 5.9. For the sake of simplicity, we ignore future expected payoff under the assumption that the optimality of signal precisions is not impacted.

$$\rho_t^* = \arg\max_{\rho_t} \{\mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t] - c(\rho_t)\} \tag{5.10}$$

To find the optimal signal precision ρ_t^*, take the first derivative of the objective function with respect to ρ and set it equal to 0, so that

$$\frac{\partial}{\partial \rho_t}(\mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t] - c(\rho_t)) = 0$$

$$\implies \frac{\partial \mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t]}{\partial \rho_t} = \frac{c(\rho_t)}{\rho_t}$$

Given the convexity of $c(\rho_t)$, the cost function penalizes deviations from a common precision ρ_0. Therefore, from Equation 5.6

$$\implies \frac{\partial \mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t]}{\partial \rho_t} = 2\varpi(\rho_t - \rho_0)$$

This can be rearranged to solve for the optimal precision ρ_t^*,

$$\rho_t = \rho_0 + \frac{1}{2\varpi} \cdot \frac{\partial \mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t]}{\partial \rho_t} \tag{5.11}$$

Here, ϖ and ρ_0 are common across players. So if we assume that the marginal expected utility of increasing precision $\frac{\partial \mathbb{E}[U(x_{it}, \theta, \tilde{x}^t)|\rho_t]}{\partial \rho_t}$ is the same for all players (under the assumption that all players face the same environment and have similar preferences), the optimal precision

ρ_t^* will be the same for every player. Finally, we see that the structure of the information cost function naturally leads to a common precision among all players. This is primarily because any player who deviates from the optimal precision incurs a significantly high cost, making it optimal for all players to choose the same level of precision in equilibrium. We eventually depict how this factor implies two further implications which would play a major role in obtaining uniqueness in equilibrium outcomes.

Threshold Strategy

Having finally proved that all players choose the same precision ρ_t^* for their signals in equilibrium, we can write that across each period t, a player's signal at equilibrium is given by : $x_{it} = \theta + \zeta_{it}$, where $\zeta_{it} \sim N(0, \sigma_t^{*^2})$ and $\rho_t^* = \frac{1}{\sigma_t^{*^2}}$
Let us hypothesize the players to be deciding whether to attack or not based on the signal. A player i follows a threshold strategy with cutoff ϱ_t in each period t if we have :

$$a_{it} = \begin{cases} 0(\text{Attack}) & x_{it} > \varrho_t \\ 1(\text{Not Attack}) & x_{it} \leq \varrho_t \end{cases}$$

Given the common precision ρ_t^*, the distribution of x_{it} is also identical for all players. Therefore, the threshold ϱ_t that optimizes their strategy will be identical for all players, leading to a common threshold strategy.

Homogeneous Belief Updates

As a second implication of all players receiving same precision signals at equilibrium, we show that their posterior beliefs about θ are updated in a homogeneous manner.
Simply from the method of Bayesian Updating of prior beliefs, we deduce that the posterior distribution of θ for each player after observing the signal x_{it} is given by :

$$\theta | x_{it} \sim N(\mu_{it}, \sigma_{\theta,t}^2)$$

where

$$\mu_{it} = \frac{\sigma^2_{\theta,(t-1)}\rho^*_t x_{it} + \sigma^{*^2}_t \mu_{t-1}}{\sigma^2_{\theta,(t-1)}\rho^*_t + \sigma^{*^2}_t}$$

and

$$\sigma^2_{\theta,t} = \frac{\sigma^2_{\theta,(t-1)}\sigma^{*^2}_t}{\sigma^2_{\theta,(t-1)}\rho^*_t + \sigma^{*^2}_t} = \frac{\sigma^2_{\theta,(t-1)}\sigma^{*^2}_t}{\sigma^2_{\theta,(t-1)}/\sigma^{*^2}_t + 1}$$

Clearly, since $\sigma^{*^2}_t$ is common across players, the belief update process will yield similar posterior beliefs μ_{it} and the variances $\sigma^2_{\theta,t}$ across all players.

Final Result

From the discussions about the two possible implications of common equilibrium signal precision across all players in each period t, we gather an important takeaway : When all players in a dynamic global game share a common signal precision, they naturally converge on using the same threshold strategy to make decisions, such as whether to attack or not. This shared threshold, derived from their homogeneous belief updates, ensures that players' actions are coordinated, as they are basing their decisions on a similar level of information (through identical precisions) and similar expectations about the future. Therefore, uncertainty is reduced, leading to a situation where players' actions are synchronized. To conclude, a unique threshold equilibrium persists under the assumption that there is a common signal precision. More formally, the two conditions for a unique equilibrium are as follows, obviously under the hypothesis of common signal precision :

- Threshold strategy : Probability of attacking is identical across all players.

$$\mathbb{P}(\text{ Player } i \text{ attacks}|\mu_{t-1}, \rho^*_t, \sigma^2_{\theta,t}) = \mathbb{P}(x_{it} > \varrho_t|\mu_{t-1}, \rho^*_t, \sigma^2_{\theta,t})$$

- Equilibrium is unique if and only if ϱ_t is uniquely determined by common posterior μ_{it}. Unique posterior further leads to a unique ϱ_t

Chapter 6

Conclusion

The novel dynamic model developed in the last chapter offers crucial insights into the role of costly information acquisition in real-world crises such as bank runs and liquidity crises, where the flow of incomplete information over time plays a critical role. In this final chapter, we reflect on these learnings and their broader implications.

Our model demonstrates that when players incur cognitive costs to acquire information about fundamentals over multiple periods, and the cost structure is convex, they tend to converge on acquiring signals with the same precision. This uniformity in information acquisition leads to similar threshold strategies among players, a pattern observable in scenarios like financial markets, international relations, and technology adoption. The convergence in information precision suggests that, despite diverse motivations, market participants often develop similar beliefs about underlying fundamentals (θ) resulting in synchronized actions such as market crashes or widespread technology adoption. Furthermore, homogeneous belief updates reinforce the predictability and stability of outcomes in global games under certain conditions. Importantly, the model captures the dynamic nature of decision-making, where evolving beliefs and strategies, based on gathered information, can stabilize or destabilize the system over time. In contexts such as economic policy-making, military planning, or long-term investments, the model suggests that the path of future strategies is significantly influenced by how information is acquired and processed. Early decisions can set a precedent, leading to path-dependent outcomes. For instance, early signals about inflation or growth can have a lasting impact on interest rates or investment decisions through self-reinforcing feedback loops.

By establishing conditions for a unique equilibrium, the model empowers stakeholders to guide the underlying system toward more desirable outcomes through informed, anticipatory decisions. Policymakers can influence global games by shaping the information environment—whether by incentivizing information acquisition, enhancing transparency, or regulating information dissemination. Ensuring that players operate with a similar knowledge base can lead to more stable and predictable outcomes. Conversely, if information remains costly or unevenly distributed, it may lead to disparate actions and multiple equilibria, some of which could result in market crashes or political instability. As discussed earlier, the establishment of a unique, robust equilibrium outcome in global games is vital for ensuring systemic stability.

Extensions

Possible extensions to the field of costly information acquisition could involve incorporating correlated information signals, as discussed by Denti (2023). Introducing correlation among signals could capture situations where players' information is interdependent, reflecting more realistic scenarios where market participants or decision-makers do not operate in isolation but are influenced by shared sources or external events. Additionally, the extensive use of stochastic choice rules, as explored by Morris and Yang (2022), could add another layer of realism by accounting for the inherent randomness and noise in decision-making processes. Furthermore, the model could be enhanced by assigning additional properties to information costs, such as non-linearities that reflect diminishing returns or externalities that lead to herding behavior. Also, we casually assume that the marginal expected utility of increasing precision is constant for all players, which may not be a realistically sound hypothesis. By addressing these issues and incorporating the aforementioned elements, the model could better capture complex dynamics observed in real-world scenarios, such as the tendency of individuals to follow the crowd or overreact to shared information, thereby offering richer insights into the implications of costly information acquisition.

References

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan (2004). "Information Dynamics and Equilibrium Multiplicity in Global Games of Regime Change". *NBER Working Paper* w11017.

— (2006). "Signaling in a Global Game: Coordination and Policy Traps". *Journal of Political economy* 114.3, pp. 452–484.

— (2007). "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks". *Econometrica* 75.3, pp. 711–756.

Angeletos, George-Marios and Alessandro Pavan (2007). "Efficient Use of Information and Social Value of Information". *Econometrica* 75.4, pp. 1103–1142.

Angeletos, George-Marios and Ivan Werning (2006). "Crises and Prices: Information Aggregation, Multiplicity, and Volatility". *American Economic Review* 96.5, pp. 1720–1736.

Atkeson, Andrew (2000). "Rethinking Multiple Equilibria in Macroeconomic Modelling". *NBER Macroeconomics Annual* 15, pp. 162–171.

Bergemann, Dirk and Stephen Morris (2016). "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games". *Theoretical Economics* 11.2, pp. 487–522.

Blackwell, David (1953). "Equivalent Comparisons of Experiments". *The Annuals of Mathematical Statistics*, pp. 265–272.

Calvo, Guillermo A (1998). "Capital Flows and Capital-Market Crises: The Simple Economics of Sudden Stops". *Journal of Applied Economics* 1.1, pp. 35–54.

Carlsson, Hans and Eric Van Damme (1993). "Global Games and Equilibrium Selection". *Econometrica: Journal of the Econometric Society*, pp. 989–1018.

Cole, Harold L and Timothy J Kehoe (1996). "A Self-Fulfilling Model of Mexico's 1994–1995 Debt Crisis". *Journal of International Economics* 41.3-4, pp. 309–330.

Dasgupta, Amil (2007). "Coordination and Delay in Global Games". *Journal of Economic Theory* 134.1, pp. 195–225.

Dean, Mark and Nathaniel Neligh (2023). "Experimental Tests of Rational Inattention". *Journal of Political Economy* 131.12, pp. 3415–3461.

Denti, Tommaso (2023). "Unrestricted Information Acquisition". *Theoretical Economics* 18.3, pp. 1101–1140.

Fosgerau, Mogens and Gege Jiang (2019). "Travel Time Variability and Rational Inattention". *Transportation Research Part B: Methodological* 120, pp. 1–14.

Gilli, Mario (2001). "A General Approach to Rational Learning in Games". *Bulletin of Economic Research* 53.4, pp. 275–303.

Goldstein, Itay and Ady Pauzner (2005). "Demand Deposit Contracts and the Probability of Bank Runs". *The Journal of Finance* 60.3, pp. 1293–1327.

Harsanyi, John C (1973). "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points". *International Journal of Game Theory* 2.1, pp. 1–23.

— (1968). "Part II: Bayesian Equilibrium Points". *Management Science* 14, pp. 320–334.

Hartigan, JA (1983). "Asymptotic Normality of Posterior Distributions". *Bayes Theory*, pp. 107–118.

Hébert, Benjamin and Michael Woodford (2023). "Rational Inattention when Decisions Take Time". *Journal of Economic Theory* 208, p. 105612.

Hoshino, Tetsuya (2018). *Rational Inattention in Games. Technical Report, Pennsylvania State University.*

Latsis, John (2009). "Hume and The Concept of Convention". *Recherches sur la Philosophie et le Langage* 26, pp. 217–34.

Maćkowiak, Bartosz and Mirko Wiederholt (2009). "Optimal Sticky Prices under Rational Inattention". *American Economic Review* 99.3, pp. 769–803.

Matějka, Filip and Alisdair McKay (2015). "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model". *American Economic Review* 105.1, pp. 272–298.

Mertens, Jean -Francois and Shmuel Zamir (1985). "Formulation of Bayesian Analysis for Games with Incomplete Information". *International Journal of Game Theory* 14, pp. 1–29.

Milgrom, Paul and John Roberts (1990). "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities". *Econometrica: Journal of the Econometric Society*, pp. 1255–1277.

Morris, Stephen and Hyun Song Shin (1998). "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks". *American Economic Review*, pp. 587–597.

— (2000). "Rethinking multiple equilibria in macroeconomic modeling". *NBER Macroeconomics Annual* 15, pp. 139–161.

— (2002). "Social value of public information". *American Economic Review* 92.5, pp. 1521–1534.

— (2003). "Global games: Theory and applications". *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1*. Cambridge University Press, pp. 56–114.

— (2004). "Coordination Risk and the Price of Debt". *European Economic Review* 48.1, pp. 133–153.

— (2006). "Inertia of Forward-looking Expectations". *American Economic Review* 96.2, pp. 152–157.

Morris, Stephen and Ming Yang (2022). "Coordination and Continuous Stochastic Choice". *The Review of Economic Studies* 89.5, pp. 2687–2722.

Obstfeld, Maurice (1986). "How Integrated are World Capital Markets? Some New Tests". *NBER Working Paper* w2075.

— (1996). "Models of Currency Crises with Self-Fulfilling Features". *European Economic Review* 40.3-5, pp. 1037–1047.

Rubinstein, Ariel (1989). “The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"”. *The American Economic Review*, pp. 385–391.

Sims, Christopher A (2003). “Implications of Rational Inattention”. *Journal of Monetary Economics* 50.3, pp. 665–690.

Szkup, Michal and Isabel Trevino (2015). “Information Acquisition in Global Games of Regime Change”. *Journal of Economic Theory* 160, pp. 387–428.

Yang, Ming (2015). “Coordination with Flexible Information Acquisition”. *Journal of Economic Theory* 158, pp. 721–738.

Yildiz, Muhamet and Jonathan Weinsten (2004). “Impact of Higher-Order Uncertainty”. *Econometric Society 2004 North American Winter Meetings*. 157. Econometric Society.