Econometrica, Vol. 88, No. 1 (January, 2020), 297–335

INFORMATIONAL CHANNELS OF FINANCIAL CONTAGION

ISABEL TREVINO
Department of Economics, University of California, San Diego

Two main classes of channels are studied as informational sources of financial contagion. One is a fundamental channel that is based on real and financial links between economies, and the second is a social learning channel that arises when agents base their decisions on noisy observations about the actions of others in foreign markets. Using global games, I present a two-country model of financial contagion in which both channels can operate and I test its predictions experimentally. The experimental results show that subjects do not extract information optimally, which leads to two systematic biases that affect these channels directly. Base-rate neglect leads subjects to underweight their prior, and thus weakens the fundamental channel. An overreaction bias strengthens the social learning channel, since subjects rely on information about the behavior of others, even when this information is irrelevant. These results have significant welfare effects rooted in the specific way in which these biases alter behavior.

KEYWORDS: Contagion, global games, experiments, social learning, behavioral biases.

1. INTRODUCTION

FINANCIAL GLOBALIZATION HAS GIVEN rise to a large number of financial crises that are rapidly transmitted across countries, a phenomenon known as financial contagion (see Schmukler, Zoido-Lobaton, and Halac (2006)). Different authors have emphasized different channels for the propagation of crises through contagion. Among the plethora of channels studied in the literature, we can distinguish two main classes: one based on fundamental links and the other based on social learning.[1] Contagion occurs through the fundamental channel when a crisis spreads across countries because of common shocks that affect their fundamentals, for example, through trade or the financial sector. Contagion can also occur between countries with weak or no fundamental links that share external characteristics, which lead investors to fear a crisis in one country after observing a crisis in a similar market. According to this view, contagion can arise as a result of social imitation when investors base their behavior on noisy observations of the actions of others in foreign markets. In this case, investors act according to their beliefs about the apparent similarity between the two markets, and their beliefs become self-fulfilling when the crisis (which could have been avoided) is spread to the second country. In both of these cases, speculation is exacerbated by incomplete and asymmetric information among investors

Isabel Trevino: itrevino@ucsd.edu

I would like to thank seven anonymous referees for insightful comments and suggestions, as well as Andrew Schotter, Jess Benhabib, Peter Bossaerts, John Duffy, Ignacio Esponda, Guillaume Frechette, Douglas Gale, Antonio Guarino, Bernard Herskovic, Tatiana Kornienko, Nic Kozeniauskas, Alessandro Lizzeri, Adam Meirowitz, Stephen Morris, Ryan Oprea, David Pearce, Debraj Ray, Joel Sobel, Charlie Sprenger, Michal Szkup, Sevgi Yuksel, and seminar participants at various institutions and conferences. I would like to thank the National Science Foundation via grant SES-105962, the Center for Experimental Social Science at NYU, and the Economics department at UCSD for financial support. All remaining errors are mine.

[1]Claessens and Forbes (2001) compile a series of papers that study the contagious episodes in the 1990s to understand the different channels that could have been responsible for contagion. Kaminsky, Reinhart, and Vegh (2003) emphasize the importance of the fundamental and social learning channels. One example of a paper that studies financial links is Kodres and Pritsker (2002), and an example of a model of contagion through social learning is Calvo and Mendoza (2000).

 https://doi.org/10.3982/ECTA15604

and between investors and governments. However, the mechanisms that lead to a crisis in these two scenarios are different. Taking debt crisis contagion as an example, in the first case a country might default because it is unable to honor its debt on account of insolvency that is related to the insolvency of another country that has defaulted. In the latter case, a default might occur as a result of the illiquidity caused by a mass withdrawal of funds based on speculation after observing agents withdraw their funds in a similar market.[2]

This paper addresses, both theoretically and experimentally, the effects and interaction of fundamental and social information in the context of financial contagion. I develop a theoretical model of financial contagion based on global games, and then test it experimentally to provide some empirical evidence as to how these channels operate. The theoretical model provides the framework for the experiment and has desirable features, such as the tension between strategic and fundamental uncertainty, while providing enough simplicity to isolate the strengths of the two channels via separate parameters. The purpose of the experimental analysis is not just to test a theoretical model, but to characterize how public information related to fundamentals and to social behavior affects individual decisions in different environments and how this can lead to contagion.

As pointed out by Goldstein (2012), differentiating between these channels is crucial for policy analysis. However, there is no conclusive empirical evidence on the matter. Although a large empirical literature has established a strong link between contagious episodes and fundamentals (see, for example, Caramazza, Ricci, and Salgado (2004) and Kaminsky, Lyons, and Schmukler (2004)), not all episodes of contagion can be explained by fundamental links. Some studies focus on the role of panics due to social imitation in generating contagion, but the social learning channel is difficult to test empirically because it depends crucially on the information sets of agents, which are typically unobservable.[3] The laboratory is a useful methodological tool for investigating this type of question, as it provides an ideal environment for studying the reactions of agents to different types of information, while controlling exogenously the strength of fundamental links and the accuracy of social information. Clearly, an experimental session cannot recreate exactly the decisions that investors face in financial markets. However, the tensions and trade-offs are qualitatively mirrored, so that we can interpret the behavior of experimental subjects as a qualitative guide to the type of behavior that financial market participants might exhibit.

The model of financial contagion presented in this paper has global games as a building block. Global games are coordination games with incomplete information where agents do not know the underlying state of the economy, which determines their payoffs. Instead, they receive noisy private and public signals about it and have to make inferences regarding the realization of the state and the likely actions of others. This perturbation in the information structure, first introduced by Carlsson and van Damme (1993), leads to a

[2]These two channels, however, are not mutually exclusive. For example, it might be rational for agents to follow the actions of others in a foreign country as long as the two countries are linked through fundamentals. However, it is also possible that social imitation serves as a channel of contagion in the absence of fundamental links when agents panic after seeing the actions of unrelated others.

[3]Some studies look for evidence for the social learning channel of contagion. Kaminsky and Schmukler (1999), for example, studied the type of news that triggered stock price fluctuations in the Asian markets in 1997–1998. They suggest that herding behavior was responsible for the changes that cannot be explained by any apparent substantial news. With a similar database, Basu (2002) interpreted residual persistence in spillover of shocks after controlling for fundamentals to be indicative of learning-based contagion. These papers point out successfully that not all cases of contagion can be explained through fundamentals. However, this type of inference does not provide conclusive evidence for the social learning channel.

very rich architecture of higher-order beliefs and ultimately selects a unique equilibrium in threshold strategies. This feature contrasts with earlier results on coordination games that predicted multiple equilibria, and it renders global games very suitable for policy analysis by focusing only on one possible outcome.[4] The use of global games to model financial contagion was initially studied by Dasgupta (2004) and Goldstein and Pauzner (2004), who focus on only one channel of contagion (capital connections of banks in the former, investors' wealth in the latter) to demonstrate that contagion can be an equilibrium outcome in a global game.[5]

In this paper, I apply the techniques of global games to financial contagion in a model with two economies whose fundamentals are correlated and who are vulnerable to runs on the funds used to finance their debt. The speculative run in each country is modeled as a global game in which investors in each country receive noisy private signals about the state of the economy and have to decide whether to withdraw their funds or to roll over their loans until maturity. The model is sequential: In the first period, agents in the first country make their decisions based on their prior information and on private signals about the state of the economy. In the second period, agents in the second country know the level of correlation between the two fundamentals and the prior about the state in Country 1 (fundamental link), they receive a noisy public signal about the proportion of agents that withdraw their funds in the first country (social learning link), and they also observe an informative private signal about the fundamentals in their own country.[6] Using global games as the workhorse for this model has two main benefits. First, it preserves the fundamental and strategic uncertainty inherent in the speculative episode in each country. Second, it provides a simple way to keep track of, and vary experimentally, the strengths of the two channels of contagion with the use of only two parameters: the correlation between fundamentals and the precision of the signal about the behavior of agents in the first country. The theoretical model illustrates the importance of prior beliefs in determining the direction of comparative statics with respect to these two parameters. I design ten experimental treatments that vary the strengths of the channels of contagion and the prior about the state in Country 1 in order to create environments where the information related to each of the two channels might or might not be relevant.

The experimental analysis revolves around two main hypotheses. The first one relates to the use of available information. Agents in the second country have access to three different sources of information: the prior distribution of the state in the first country (publicly observed), the public signal about the behavior of agents in the first country, and a private signal about the realized state in the second country (which is always informative). The first main hypothesis is that subjects would take into consideration the prior distribution of the state in the first country only when the states are correlated, in which case we say that the fundamental channel is active, and that they would take into account the signal about the behavior of agents in the first country (i.e., the social learning channel would be active) only if this signal is informative, meaning that it is correlated

[4]For an overview of global games, see Morris and Shin (2003). For applications and extensions of global games models see, among others, Morris and Shin (1998, 2004), Angeletos, Hellwig, and Pavan (2006, 2007), Angeletos and Werning (2006), Hellwig, Mukherji, and Tsyvinski (2006), or Edmond (2013).

[5]Other theoretical papers that study financial contagion in a global games context are Manz (2010), Oh (2013), and Ahnert and Bertsch (2018).

[6]In order to model situations of rapid speculation, agents in Country 2 do not observe the outcome of Country 1 before making their decision; instead, they observe the prior about the state in Country 1 and a noisy signal about the actions of Country 1 agents.

to the true behavior of agents in the first country and the states of the two countries are correlated.

The experimental results show two systematic biases in information processing that affect the channels of contagion directly: base-rate neglect, where subjects underweight or ignore the prior when the correlation between states is high, weakening the fundamental channel, and an overreaction bias, where subjects systematically take into account the signal about the behavior of agents in the first country, even when this signal is completely uninformative, which strengthens the social learning channel when it should be inactive. While similar biases have previously been documented in the behavioral and experimental literature on individual decision making, the results in this paper show evidence for the emergence of these biases in a strategic context.[7] More importantly, these biases directly affect the specific mechanisms of these two channels of contagion. Therefore, these results tie well-known findings in behavioral economics to macroeconomic phenomena and point to new information-based mechanisms for the transmission of crises that have not been studied before.

The second main hypothesis relates to welfare. I study the effects of the two aforementioned biases on outcomes and payoffs and find that the way in which subjects overreact to social information has significant consequences for welfare. In most cases, they overreact to the signal about the behavior of others by following panics (withdrawing after observing a signal that others withdrew when equilibrium prescribes to roll over), even when this signal is uninformative. These contagious panics lead to welfare losses in the form of higher frequencies of withdrawals and lower realized payoffs than those prescribed by equilibrium. However, in two cases, subjects overreact to social information by following unrelated actions that show confidence (rolling over after observing a signal that others rolled over when equilibrium prescribes to withdraw). In these cases, social information serves as a coordination device and the resulting contagious confidence leads to welfare gains.

Additional analysis emphasizes how different sources of public information not only lead to different biases but they also determine their extent and persistence. The base-rate neglect bias weakens as subjects gain more experience by playing the game for a larger number of rounds, but the social imitation bias remains strong and significant, even with increased experience. This implies that overreaction to social information is persistent over time and that this type of information is processed differently from public information that does not reflect behavior. Moreover, subjects do not exhibit base-rate neglect when the social signal is presented as a standard public signal, that is, when it is not framed to reflect the behavior of others. Therefore, the emergence of the base-rate neglect bias seems to depend on the nature of the other available signals.[8] This has important policy implications that hint at a lack of substitutability between government announcements (standard public signals) and market observations (public signals about behavior). In models with Bayesian agents, these two types of public observations should not lead to different predictions, because public information about behavior is

[7]In particular, the overreaction to social information is intrinsically different from overreaction in individual decision-making environments, for two main reasons: The information it conveys is rooted in social interaction, and, by virtue of being a publicly observed signal, it might serve as a coordination device.

[8]A possible interpretation of this result is that processing social information might be cognitively more costly in terms of the effort needed by subjects to understand the reasoning of the people whose behavior is summarized by the signal. Given this higher cognitive effort, subjects are more prone to ignore other available information, leading to base-rate neglect.

observationally equivalent to a standard public signal. However, the results of this experiment show that people respond differently to public information that reflects human behavior in a persistent manner and that this overreaction to social information can either amplify or dampen the effects of contagion. The differentiated use of public information documented in this paper implies that the causes and risk factors for contagion might be different in severity, source, and solution than suggested by conventional models.

The analysis presented in this paper is done in the context of financial contagion; however, the experiment was conducted using a neutral frame, so the results could easily be applied to other relevant contexts. One example is bank-run contagion, where the fundamental link could be interpreted as a common exposure of two banks to asset shocks or macroeconomic shocks that affect the liquidity in both banks. Excessive withdrawals in one bank could lead depositors in other banks to withdraw their money if they fear that their bank might be insolvent in the future, and thus lead to a liquidity crisis in the second bank as a result of social contagion.[9] Another context where fundamentals and social learning play a key role in contagion is political uprisings, such as the contagious episodes of political regime change that took place during the Arab Spring in 2010–2012. In this case, citizens decide whether or not to protest the repressive government, in the hope of starting a revolution that will overthrow the regime. The fundamental link can be thought of as how similar (radical) the ideologies of the governments are, and the social information corresponds to the number of people that have taken to the streets in protest in other countries.[10]

The paper is structured as follows. Section 2 presents the theoretical global games model of financial contagion. Section 3 presents the experimental design. The experimental results are presented in Section 4. Section 5 relates the findings of the paper to the existing literature. Section 6 discusses the implications of the results and concludes.

2. THEORETICAL MODEL

I first introduce a model with a continuum of agents in each country, continuous state and signal spaces, and continuous probability distributions. This serves as a formal characterization of the environment to illustrate the main theoretical implications in terms of the two channels of contagion. The full characterization of this model is relegated to an appendix on the author's website.[11] I then discretize the setup to derive predictions for the experiment. Since this is a fairly complex model, the setup is discretized in every possible dimension in order to make it more accessible to the experimental subjects. The discretized version mirrors the qualitative properties of the formal model for the parameters used in the experiment. The general setup is presented below, followed by two subsections that spell out the specific features of the continuous and discrete models.

There are two countries in the economy, Country 1 and Country 2, and a different set of agents in each country. There are two periods, and agents related to Country $n \in \{1, 2\}$ are active only in period n. For simplicity, countries become active in the order of their numeraire.

[9]For example, the run on the British bank Northern Rock in 2007 led to fears of contagion across the banking system in the UK. The Bank of England responded by providing liquidity support to Northern Rock in order to contain the spread of a deposit run.

[10]Starting with the Tunisian revolution in 2010, by the end of 2012 leaders had been forced out of power in Tunisia, Egypt, Libya, and Yemen, with uprisings also in neighboring countries such as Bahrain and Syria.

[11]See http://econweb.ucsd.edu/~itrevino/pdfs/continuum_model.pdf.

Both countries use standard debt contracts to finance their debt. These contracts specify an interim stage where agents (creditors) can review their investment and decide whether to roll over their loan to maturity or to withdraw their funds prematurely. Formally, agent i in Country n chooses an action $a_n^i \in \{0, 1\} = \{\text{withdraw, roll over}\}$. Agents from Country n have funds invested only in Country n.

The setup in each country follows the one-country setup of Morris and Shin (2004), where domestic agents buy securities to finance the country's government debt. Debt contracts specify two different face values, depending on the time of liquidation.[12] The face value of repayment at maturity for the debt contract in Country n is γ_n, and each agent who rolls over her loan receives this amount if the country remains solvent. If the country defaults, then agents who roll over get 0. If agents choose to withdraw their funds prematurely, they get the lower face value of early withdrawal, $\lambda_n \in (0, \gamma_n)$ with certainty. Agents might want to withdraw their funds at the interim stage if they fear that the country may default and not repay its debt. These fears are self-fulfilling, since the greater the number of agents that withdraw simultaneously, the greater the likelihood of default. Therefore, withdrawing is a safe action that leads to a constant payoff, regardless of the outcome in the country, and rolling over is a risky action that, depending on the outcome, can lead to a higher or lower payoff than withdrawing. The payoffs of an agent in Country n are summarized below:

	Solvency at maturity	Default
Roll over loan	γ_n	0
Withdraw	λ_n	λ_n

Whether Country n honors its debt at maturity or defaults depends on two factors: the underlying state of the economy (not known to agents), represented by the variable $\theta_n \in \mathbb{R}$, and the proportion of agents that withdraw in that country.

There exist $\{\underline{\theta}_n, \overline{\theta}_n\} \in \mathbb{R}$ such that $\underline{\theta}_n < \overline{\theta}_n$. Under complete information, we can distinguish three regions for the fundamentals: If $\theta_n \geq \overline{\theta}_n$, agents find it optimal to roll over their debt, irrespective of the actions of others (rolling over always yields the high face value $\gamma_n > \lambda_n$). If $\theta_n < \underline{\theta}_n$, it is optimal to withdraw the funds at the interim stage (the country always defaults, and rolling over the funds would lead to a payoff of $0 < \lambda_n$). For $\theta_n \in [\underline{\theta}_n, \overline{\theta}_n)$, there is a coordination problem where the optimal action depends on the beliefs about θ_n and about the actions of the other agents. In this model, however, agents do not observe θ_n directly but receive noisy private and public signals about it.

The two countries are linked through fundamentals, that is, θ_1 and θ_2 are correlated. A high level of fundamental comovement between these economies means poor fundamentals in one country imply poor fundamentals in the other one. This would increase the probability of default in the second country, irrespective of information.

In Country 1, agents know the prior distribution of θ_1 and, once the state is realized, they observe one noisy private signal about it. Agents in Country 2 do not observe the realization of θ_1, but they know the prior distribution of θ_1 and the correlation between θ_1 and θ_2, and they observe a noisy private signal about θ_2 and a public signal about the proportion of agents that withdraw in Country 1. The latter signal incorporates a component of social learning that is not present in the standard model of global games.

[12]Two different face values for short- and long-term debt are also studied in Szkup (2015). However, in that model there is no possibility for contagion.

Theoretical model is built to explicitly highlight the differences between rational play and departures from rationality that might arise in the use of information in Country 2. As will be shown for both the continuous and discrete models, a necessary condition for agents in Country 2 to take into account the signal about the behavior of agents in Country 1 is that the fundamentals in the two countries are correlated. That is, the only scenario in which a rational player would follow the actions of others is when these actions convey information about his own payoff-relevant state, which requires fundamental correlation. This implies that, in the model, the social learning channel cannot operate independently of the fundamental channel. One advantage of setting up the model in this way is that we will be able to label as overreaction to social information any instance where experimental subjects take into account the social signal when, according to the rational benchmark, they should not, without having to rely on magnitudes of effects, which could lead to external validity concerns (see Section 3.2.1).

Since payoffs in each country do not depend on the actions taken by agents in the other country, there are no strategic considerations across periods. Therefore, the problem is simplified to a sequence of two static global games where the state in the first game affects the realization of the state in the second one.

Sections 2.1 and 2.2 characterize the formal model with a continuum of agents and the discretized model that is implemented experimentally, respectively.

2.1. *Model With a Continuum of Agents in Each Country*

There is a continuum of agents in each country, indexed by $i_n \in [0, 1]$, $n = 1, 2$. The face value of repayment at maturity, γ_n, is set to 1. Whether Country n honors its debt at maturity or defaults is determined by comparing the realization of θ_n to the proportion of withdrawing agents, l_n, as follows:

$$\text{Country } n \begin{cases} \text{remains solvent} & \text{if } l_n \leq \theta_n, \\ \text{defaults} & \text{if } l_n > \theta_n. \end{cases}$$

We can think of θ_n as fundamentals that reflect the ability of the government to meet short-term claims from agents (e.g., as an index of liquidity). Given this characterization, $\underline{\theta}_1 = \underline{\theta}_2 = 0$ and $\overline{\theta}_1 = \overline{\theta}_2 = 1$. That is, under complete information agents roll over if $\theta_n \geq 1$, they withdraw their funds if $\theta_n < 0$, and the coordination region corresponds to values of $\theta_n \in [0, 1)$.

The value of θ_1 is drawn from a normal distribution with mean μ and precision τ, $\theta_1 \sim N(\mu, \tau^{-1})$. This is common knowledge to all agents in both countries. Since events in Country 2 take place after events in Country 1 have occurred, θ_2 depends on θ_1 in that the realization of θ_1 is set to be the mean of the distribution from which θ_2 is drawn, that is, $\theta_2 \sim N(\theta_1, \tau_s^{-1})$. The precision parameter τ_s captures the link between fundamentals between the two countries.[13] Agents in Country 2 know how θ_2 depends on θ_1, but they do not observe the realization of θ_1.

Since there are no strategic considerations across periods, I solve the two subgames separately and then address the effects of the outcome in Country 1 on the outcome in Country 2.

[13]Even though this is not strictly a measure of correlation, it has the same interpretation, since an increase in τ_s increases the probability that the realization of θ_2 is closer to that of θ_1. The correlation between fundamentals was modeled in this way to reflect the sequential nature of the game.

Besides holding prior beliefs about θ_1, agents in Country 1 also observe noisy private signals about their payoff-relevant state, θ_1, given by

$$x_1^i \sim N\big(\theta_1, \tau_{r_1}^{-1}\big),$$

where the x_1^i are *i.i.d.* across $i \in [0, 1]$.

Notice that the game in Country 1 corresponds to a standard static global game, so a unique equilibrium in monotone strategies is guaranteed as long as private signals are precise enough with respect to the prior, that is, $\frac{\sqrt{\tau_{r_1}}}{\tau} > \frac{1}{\sqrt{2\pi}}$.[14] This unique equilibrium is characterized by a pair of thresholds $\{\theta_1^*, x_1^*\}$ such that agents in Country 1 withdraw their funds if they observe a low private signal $x_1^i < x_1^*$, and they roll over if $x_1^i \geq x_1^*$. Likewise, Country 1 defaults if the state of fundamentals is low, $\theta_1 < \theta_1^*$, and it stays solvent and honors its debt if the state is sufficiently high, $\theta_1 \geq \theta_1^*$. These thresholds depend on the parameters of the model.

In Country 2, the information structure is more complex, but the equilibrium follows the same logic as in Country 1. Agents in Country 2 observe private signals about the state in their *own* country, θ_2, given by $x_2^i \sim N(\theta_2, \tau_{r_2}^{-1})$, where the x_2^i are *i.i.d.* In addition, agents in Country 2 observe a public signal about the actions of agents in Country 1, given by

$$y|\theta_1 \sim N\big(\Phi^{-1}(l_1), \tau_\alpha^{-1}\big),$$

where $l_1 = \Pr(x_1^i < x_1^*) = \Phi(\frac{x_1^* - \theta_1}{\tau_{r_1}^{-1/2}})$ is the proportion of agents in Country 1 that withdraw their funds.[15] Notice that since $y \sim N(\frac{x_1^* - \theta_1}{\tau_{r_1}^{-1/2}}, \tau_\alpha^{-1})$, y is theoretically equivalent to a standard public signal about θ_1.

Agents in Country 2 know that the realization of their payoff-relevant state is determined by $\theta_2 \sim N(\theta_1, \tau_s^{-1})$. However, they do not observe the realization of θ_1. They form beliefs about the mean of this distribution by combining the prior about θ_1 (determined by μ and τ) and the observation of y using Bayes' rule. That is, for them $\theta_2 \sim N(E(\theta_1|y), \tau_s^{-1})$.[16]

Once θ_2 is realized, agents in Country 2 observe their private signals about it, $x_2^i \sim N(\theta_2, \tau_{r_2}^{-1})$. Taking into consideration all the information at their disposal (the prior about θ_1, the public signal about the behavior of agents in Country 1, and their private signals), agents in Country 2 form posterior beliefs about θ_2. As in Country 1, there is a unique equilibrium characterized by a pair of thresholds $\{\theta_2^*, x_2^*\}$ such that agents in Country 2 withdraw their funds if they observe a low private signal $x_2^i < x_2^*$, and they roll them over otherwise. Also, Country 2 defaults if $\theta_2 < \theta_2^*$, and it remains solvent if $\theta_2 \geq \theta_2^*$. The thresholds $\{\theta_2^*, x_2^*\}$ depend crucially on the parameters of the model in Country 1 and Country 2. A sufficient condition for a unique equilibrium in Country 2 is that private signals be informative enough with respect to public information, that is, $\frac{\sqrt{\tau_{r_2}}}{(\tau_s^{-1} + (\tau + \tau_{r_1}\tau_\alpha)^{-1})^{-1}} > \frac{1}{\sqrt{2\pi}}$.

[14] See Morris and Shin (2003), Hellwig (2002), or Morris and Shin (2004) for details.

[15] This transformation assumes monotonic strategies by agents in Country 1, so we restrict attention to strategies of this type. The transformation facilitates the analysis and follows Dasgupta (2007).

[16] A detailed derivation of the model is available on the author's website, http://econweb.ucsd.edu/~itrevino/pdfs/continuum_model.pdf. Adding a signal about the realization of the state in Country 1 or the outcome in Country 1 would be an interesting extension.

Notice that the activity in Country 1 affects the beliefs of agents in Country 2 through two channels: (1) the correlation between fundamentals, τ_s, which determines how relevant it is for agents in Country 2 to pay attention to the information related to Country 1. This measures the strength of the fundamental link between countries. (2) The signal about the actions of agents in Country 1, y, and its precision, τ_α, for a given level of fundamental correlation. A higher proportion of agents that withdraw in Country 1 leads agents in Country 2 to believe that fundamentals in Country 2 are weak if the states are highly correlated, so τ_α determines the accuracy with which this information is transmitted from Country 1 and Country 2. Therefore, y and τ_α characterize the social learning channel. Notice that both of these two channels are informational channels, that is, both fundamentals and social learning can lead to contagion through the public information which is revealed to agents in Country 2.

To study comparative statics, we look at changes in the probability of default in Country 2, measured by θ_2^*, that result from changing the parameters that determine the strength of the two channels of contagion, assuming that the conditions for uniqueness of equilibrium hold.[17,18]

For the fundamental channel, we observe that (1) if the equilibrium probability of default in Country 2 is low and agents have an optimistic prior about θ_2, then a higher correlation between the two countries (τ_s) will further decrease the probability of default in Country 2, and (2) if the equilibrium probability of default in Country 2 is high and agents have a pessimistic belief about θ_2, then a higher correlation between the two countries will increase the probability of default in Country 2. This result has an intuitive interpretation: Agents being optimistic about fundamentals in Country 2 before observing their private signal effectively means that they are optimistic about the realization of θ_1. In this case, an increase in τ_s leads agents in Country 2 to assign a higher weight to these optimistic beliefs, which decreases the probability of default in Country 2. This illustrates the positive effects of the fundamental link on contagion. Analogous reasoning illustrates the negative effect of the fundamental link on contagion when agents are pessimistic about the realization of θ_1.

To analyze the social learning channel of contagion, I look at the effect on the probability of default in Country 2 which results from changes in the precision of y, τ_α. To do this, we need to decompose the notion of optimistic (pessimistic) beliefs about the state in Country 2, which correspond to the posterior about θ_1, $E(\theta_1|y)$. On the one hand, τ_α, just like τ_s, affects the probability of default because it determines the weight given to public information, so the effect of a change in τ_α depends on whether agents are optimistic or pessimistic about θ_2. However, the total effect of a change in τ_α on the probability of default in Country 2 is more complex, since a change in τ_α also determines whether agents are optimistic or pessimistic in the first place. This means that there are two effects of an increase in τ_α, which can go in opposite directions. The first effect is a "coordination effect" that enhances coordination by increasing the weight given by all agents to public

[17]In the one-country setup, Morris and Shin (2004) showed that the mean of the prior has important effects on the probability of default. A country is able to stay solvent for a wider range of fundamentals (lower θ_1^*) when creditors hold an optimistic prior about the state of the economy (higher μ).

[18]The derivations of these and other comparative statics, such as those related to the precisions of the private signals, τ_{r_n}, the mean of the prior in Country 1, μ, the precision of the prior for Country 1, τ, and the payoffs of early withdrawal, λ_n, for both countries are available on the author's website, http://econweb.ucsd.edu/~itrevino/pdfs/continuum_model.pdf.

information about θ_2 (similar to the effect of an increase in τ_s). The second is an "information effect" that determines the type of beliefs that agents hold about θ_2, that is, whether they are optimistic or pessimistic about it.[19]

The main takeaway of this analysis for the experimental design is that prior beliefs about θ_2 are crucial determinants of the way in which the probability of default in Country 2 is affected by changes in the strengths of the two channels of contagion. Below we show that this tension is also present in the discrete setup with the use of numerical simulations.

2.2. *Discrete Setup*

This section introduces the model that is tested experimentally. This simplified model has two agents and discrete state and signal spaces, to make the setup easier to understand for the experimental subjects while keeping enough structure to carry over the main tensions and trade-offs of the continuous setup.[20] Even if we lose some desirable features that allow for an analytical characterization of the continuous model, numerical simulations show that the qualitative predictions translate to the discrete setup.

There are two players in each country. The state θ_n can be either low, medium, or high, $\theta_n \in \{L_n, M_n, H_n\}$. To maintain the structure of a global game, I assume that when the state is high ($\theta_n = H_n$) the country always remains solvent (rolling over is a dominant strategy), when the state is low ($\theta_n = L_n$) the country always defaults (withdrawing is a dominant strategy), and when the state is medium ($\theta_n = M_n$) the country remains solvent only if the two agents coordinate on rolling over. The face value of repayment at maturity for the debt contract in Country n is 20, and the face value of early withdrawal is 4.[21] That is, $\gamma_1 = \gamma_2 = 20$ and $\lambda_1 = \lambda_2 = 4$.

Therefore, for each possible realization of the state in Country n, agents face one of the matrices of payoffs from Table I. The headings "R" and "W" correspond to Rollover and Withdraw, respectively.

The prior distribution of θ_1 is characterized by unconditional probabilities p and q: $\Pr(L_1) = p$, $\Pr(M_1) = q$, and $\Pr(H_1) = 1 - p - q$. This is known to agents in both countries.

TABLE I

PAYOFFS, COUNTRIES 1 AND 2

$\theta_n = L_n$	R	W
ROLLOVER	0, 0	0, 4
WITHDRAW	4, 0	4, 4

$\theta_n = M_n$	R	W
ROLLOVER	20, 20	0, 4
WITHDRAW	4, 0	4, 4

$\theta_n = H_n$	R	W
ROLLOVER	20, 20	20, 4
WITHDRAW	4, 20	4, 4

[19]For example, changes in the precision of the social signal might lead beliefs to switch from optimism to pessimism (or vice versa) by the information effect. Then the coordination effect will either increase or decrease the probability of default in Country 2, depending on whether these "updated" beliefs are optimistic or pessimistic.

[20]Some global games experiments have used continuous signal and state spaces (see Heinemann, Nagel, and Ockenfels (2004) or Szkup and Trevino (2019)). However, the information structure for agents in Country 2 is significantly more complex than in these papers, and the richness of continuous state and signal spaces and probability distributions could have led to a cognitive overload for experimental subjects.

[21]These payoffs correspond to the values (in dollars) used in the experiment.

TABLE II

CONDITIONAL DISTRIBUTION OF PRIVATE SIGNALS, COUNTRY 1

	L_1	M_1	H_1
$\Pr(l_1\|\cdot)$	r	$\frac{(1-r)}{2}$	$\frac{(1-r)}{2}$
$\Pr(m_1\|\cdot)$	$\frac{(1-r)}{2}$	r	$\frac{(1-r)}{2}$
$\Pr(h_1\|\cdot)$	$\frac{(1-r)}{2}$	$\frac{(1-r)}{2}$	r

The agents in Country 1 observe private signals that can be either low, medium, or high, $x_1^i \in \{l_1, m_1, h_1\}$, and the conditional distribution of each private signal is characterized by a parameter r that determines its precision, as shown in Table II. I assume that $r > \frac{1}{3}$, so that the signals and states are positively correlated.

The equilibrium characterization in Country 1 depends on the specific combination of parameters. Although we cannot guarantee a unique equilibrium for a wide range of parameters with discrete state and signal spaces, the parameters used in the experiment are chosen to ensure a unique dominance-solvable equilibrium in monotone strategies, as in the continuous setup. Depending on the parameters, in the unique equilibrium agents follow one of the following monotonic strategies: always withdraw, roll over only for high signals, roll over for high and medium signals, or always roll over.

The state in Country 2 depends on the realization of the state in Country 1. The correlation between states is measured by a parameter $s \geq 1/3$, which is common knowledge to all agents. The probability distribution for the state in Country 2, $\theta_2 \in \{L_2, M_2, H_2\}$, given the realization of the state in Country 1, $\theta_1 \in \{L_1, M_1, H_1\}$, is presented in Table III.

Analogous to the continuous model, s characterizes the strength of the fundamental channel of contagion and y and α characterize the social learning channel. Note that if $s = 1$ the fundamentals are perfectly correlated, and if $s = 1/3$ there is no correlation between fundamentals, so agents in Country 2 should take into account the information related to Country 1 only when $s > 1/3$. In the experiment, therefore, s serves as a treatment variable that determines the strength of the fundamental channel of contagion.

Agents in Country 2 observe a public signal about the number of agents in Country 1 that withdraw their funds. Given that there are only two players in each country, let $w \in \{0, 1, 2\}$ be the true number of withdrawals in Country 1, and let $y \in \{0, 1, 2\}$ be the noisy signal that agents in Country 2 observe about w. I assume that, given θ_1, agents in Country 2 learn the true number of withdrawals ($y = w$) with probability α, and that they observe each of the two incorrect numbers of withdrawals with probability $\frac{(1-\alpha)}{2}$.

TABLE III

CONDITIONAL DISTRIBUTION OF THE STATE IN COUNTRY 2, GIVEN THE REALIZATION OF THE STATE IN COUNTRY 1

	L_1	M_1	H_1
$\Pr(L_2\|\cdot)$	s	$\frac{(1-s)}{2}$	$\frac{(1-s)}{2}$
$\Pr(M_2\|\cdot)$	$\frac{(1-s)}{2}$	s	$\frac{(1-s)}{2}$
$\Pr(H_2\|\cdot)$	$\frac{(1-s)}{2}$	$\frac{(1-s)}{2}$	s

TABLE IV

CONDITIONAL DISTRIBUTION OF THE PUBLIC SIGNAL ABOUT BEHAVIOR IN COUNTRY 1, FOR COUNTRY 2

	$w=0$	$w=1$	$w=2$
$\Pr(y=0\|\theta_1,\cdot)$	α	$\frac{(1-\alpha)}{2}$	$\frac{(1-\alpha)}{2}$
$\Pr(y=1\|\theta_1,\cdot)$	$\frac{(1-\alpha)}{2}$	α	$\frac{(1-\alpha)}{2}$
$\Pr(y=2\|\theta_1,\cdot)$	$\frac{(1-\alpha)}{2}$	$\frac{(1-\alpha)}{2}$	α

Therefore, $\alpha \geq 1/3$ measures the precision of this signal. If $\alpha = 1$, this signal is perfectly informative and agents in Country 2 observe exactly what agents in Country 1 did, but if $\alpha = 1/3$ this signal is completely uninformative. Table IV shows this conditional probability distribution.

Observing signal y is useful for agents in Country 2 to update their beliefs about the realization of θ_1, as long as $\alpha > 1/3$ and $s > 1/3$. This illustrates that fundamental correlation is a necessary condition for the social signal to be relevant for rational agents.

Finally, agents in Country 2 also observe a noisy private signal about θ_2, which, for simplicity, has the same structure as the private signal for agents in Country 1. The precision of this signal, $r > 1/3$, is assumed to be the same in both countries. The probability distribution of private signals $x_2^i \in \{l_2, m_2, h_2\}$, given θ_2, is analogous to that in Table II.

Taking into consideration all the information at their disposal (the prior about θ_1, the public signal y, and the private signal x_2^i), agents in Country 2 form posterior beliefs about their payoff-relevant state, θ_2. Just as in Country 1, the parameters in the experiment are chosen so that there is a unique dominance-solvable equilibrium in monotone strategies in Country 2. In this case, the actions taken by agents in Country 2 depend on private and public signals, so the monotonicity of the equilibrium strategies will be with respect to both of them. That is, for a given public signal y, the monotonicity of actions with respect to private signals $x_2^i \in \{l_2, m_2, h_2\}$ establishes a probability of withdrawing for low signals which is at least as high as that for medium signals, and similarly for medium signals compared to high signals. Likewise, for a given private signal x_2^i, the monotonicity of actions with respect to the public signal $y \in \{0, 1, 2\}$ establishes a probability of withdrawing after observing a signal that two agents in Country 1 withdraw which is at least as high as that for a signal which states that one agent withdraws, and similarly for one agent withdrawing compared to no agents withdrawing. The specific equilibrium actions, in terms of combinations of x_2^i and y, depend on the parameters and are specified in Section 3 for each treatment of the experiment.

As stated earlier, discretizing the state and signal spaces simplifies the experimental implementation at the cost of losing some convenient features of the model with continuous distributions that allow for an analytical characterization of the equilibrium and comparative statics. Since the discretized model is the one that is tested experimentally, it is important to show that the qualitative predictions of the continuous case translate to the discrete setup—in particular, the prediction that prior beliefs crucially determine how changes in the strengths of the two channels of contagion (measured by changes in s and α) affect the probability of default in Country 2.

For this purpose, we turn our attention to numerical simulations. Panel (a) of Figure 1 shows how the ex ante probability of default in Country 2 changes as the countries become more correlated, for both an optimistic prior and a pessimistic prior. Consistent with the

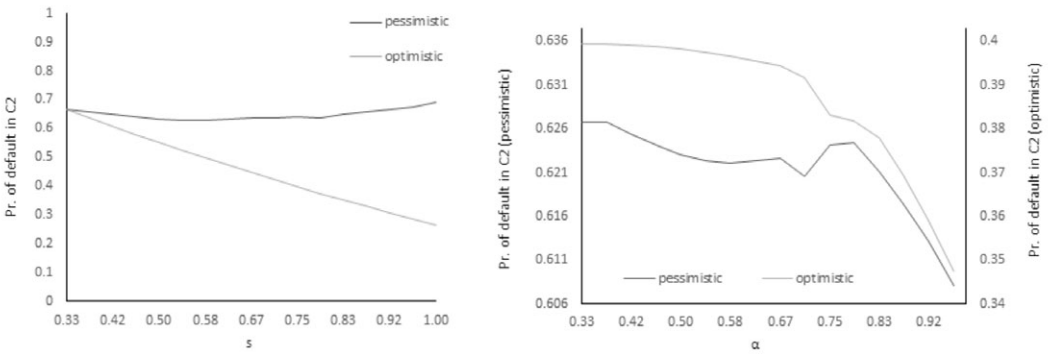

a. Correlation of fundamentals (s)

b. Precision of the social learning signal (α)

FIGURE 1.—Numerical simulations of comparative statics for the discretized model.

predictions of the continuous model, an optimistic prior leads to fewer defaults as the correlation between fundamentals increases, because agents put more weight on these optimistic beliefs. However, when the prior is pessimistic, this is not the case. Panel (b) of Figure 1 shows, on the right vertical axis, that increasing the precision of the social learning signal α leads to a monotonic decrease in the probability of withdrawals in Country 2 for the parameters used in the simulation. However, we see on the left vertical axis that when the prior is pessimistic the relationship is not monotonic.[22] Just as in the continuous model, changes in the precision of the social signal might not have a monotonic effect on the probability of default in Country 2, because of the coordination and information effects caused by varying α (given $s > 1/3$). For the parameters used in the experiment, we can see that whether these effects lead to monotone comparative statics or not also depends on the prior. These results reinforce the necessity of inducing different types of priors in the experiment.

3. EXPERIMENTAL PROCEDURES

3.1. *Parameters Used in the Experiment*

The main treatment variables in the experiment are the parameters that measure the strengths of the two channels of contagion: s and α. The combinations of these parameters in the different treatments are the following: $(s, \alpha) \in \{(1/3, 1/3), (3/4, 1/3), (1/3, 3/4), (3/4, 3/4), (1, 1)\}$. In the first three cases, at least one of the channels is inactive, because either the states are uncorrelated ($s = 1/3$), the signal about behavior of agents in Country 1 is uninformative ($\alpha = 1/3$), or both. Notice that in the third case, even if the signal about behavior of agents in Country 1 is informative ($\alpha = 3/4$), both channels are inactive since the states are uncorrelated ($s = 1/3$), so subjects in Country 2 should disregard the information from the prior about the state in Country 1 and the signal y. Given the theoretical importance of prior beliefs for the comparative statics with respect to these

[22]The parameters used in these simulations are very similar to those used in the experiment. For the optimistic prior, we set $p = 0.175$ and $q = 0.175$, while for the pessimistic prior we set $p = 0.65$ and $q = 0.175$. We set $s = 3/4$ (for panel (b)), $\alpha = 3/4$ (for panel (a)). As in the experiment, for panel (a) we set $r = 0.6$. However, for panel (b) we set $r = 0.6$ for Country 2, and $r = 0.9$ for Country 1 to get greater variability of equilibrium behavior in Country 1 when the prior is optimistic.

TABLE V

PRIOR PROBABILITY DISTRIBUTIONS

Optimistic prior			Pessimistic prior		
$\Pr(L_1)$	$\Pr(M_1)$	$\Pr(H_1)$	$\Pr(L_1)$	$\Pr(M_1)$	$\Pr(H_1)$
17.5%	17.5%	65%	65%	17.5%	17.5%

two parameters, in the experiment I induce two types of prior beliefs about the state in Country 1: one that gives a high probability to a high state in Country 1 (optimistic), and one that gives a high probability to a low state in Country 1 (pessimistic). The probability distributions corresponding to optimistic and pessimistic priors about θ_1 are given in Table V.

The precision of the private signals in each country is given by $r = 6/10$. That is, in each country the probability of observing a private signal that is consistent with the realization of the state is 6/10, and the probability of observing each of the two inconsistent signals is 2/10.

3.2. *Experimental Design*

The experiment was conducted using the usual computerized recruiting procedures. All subjects were undergraduate students. The experiment was run in two locations: at the Center for Experimental Social Science at New York University (NYU) and at the Economics Lab at the University of California, San Diego (UCSD). Most sessions lasted approximately 60 minutes, and subjects earned \$17 on average. Two longer sessions, run at UCSD, lasted approximately 90 minutes, and subjects earned \$27 on average. The experiment was programmed and conducted with the z-Tree software (Fischbacher (2007)). The instructions for the experiment can be found in the Supplementary Material (Trevino (2020)).

I implemented a between-subjects design in order to directly compare the behavior of subjects across treatments. Each session consisted of 30 independent and identical rounds (except for the longer sessions which had 50 rounds), and subjects were randomly matched in pairs in every round. In each round, subjects made decisions simultaneously without a preselected action.

For Country 1, the only treatment variable was the prior (see Table V for parameters). Since the focus of this study is to understand the behavior of agents in Country 2, the same observations from each Country 1 treatment were used as the baseline for every Country 2 session. Every combination of parameters (s, α) used in the Country 2 sessions was run twice: once where prior optimistic beliefs about Country 1 were induced, and once where prior pessimistic beliefs were induced. In every round, each pair of subjects in a Country 2 session was randomly matched with a pair of subjects from the Country 1 session with the corresponding prior beliefs. The observations made by Country 1 pairs determined both the fundamental state in Country 1 (θ_1) and the number of withdrawals in Country 1 that, depending on the parameters s and α, determined the state in Country 2 and the public signal that subjects in Country 2 received about the behavior of that specific pair of subjects in Country 1. In this way, the behavior of every pair of subjects that participated in Country 1 sessions was used as the base for one pair of subjects in each Country 2 session. Since subjects were randomly matched in pairs in every round and the matching of pairs from Country 1 to Country 2 sessions was done randomly in every round, the one-shot feature of the game was preserved. In other words, subjects in Country 2 sessions should not condition their decisions on past performance of their opponent or of subjects in Country 1, since in each round they are matched with someone new in the room and they observe information about a new pair of subjects from a Country 1 session.

Subjects in Country 2 sessions were told in the instructions the precise way in which θ_2 depended on the state in a previous experiment (Country 1 session) and the information structure of the Country 1 and Country 2 games. In each round, they observed simultaneously the public signal about the behavior of the subjects in the previous experiment and their private signal about θ_2. The induced prior, coming from the unconditional probability distribution about θ_1, is held fixed throughout the entire session and is publicly observed by both pair members, just as the social signal y. This is in contrast to the private signal, which is drawn and observed independently by each subject.

To avoid framing effects, the game was explained using neutral terms. Subjects were told to choose between two actions, A (roll over) and B (withdraw), avoiding terminology such as "withdraw," "roll over," or "default." For this reason, we can interpret the results of the experiment in settings other than financial contagion.

Each round proceeded as follows: Subjects were randomly and anonymously matched with another person in the room. Then they observed their signals (a private signal about θ_1 in Country 1 sessions; a private signal about θ_2 and a public social signal y in Country 2 sessions) and had to simultaneously make a choice between actions A and B. After each round, they received feedback about the realization of the state in their own country, the signals they observed, the outcome of the game, and their individual payoff for the round.[23]

The computer randomly selected three of the rounds played (one from rounds 1–10, one from 11–20, and one from 21–30), and subjects were paid the average of the payoffs from those rounds. All parameters in the experiment were expressed in dollar amounts.

Each of the 10 Country 2 treatments was run at both NYU and UCSD (one session in each location), and the same information from Country 1 sessions (which came from NYU) was used in both locations.

Overall, there were a total of 580 participants. Table VI summarizes the experimental design and contains the equilibrium predictions for each treatment in Country 1 and Country 2.

Treatments C2:2*, C2:7*, C2:2**, and C2:7** serve as robustness checks. C2:2* and C2:7* are identical to C2:2 and C2:7, except for the number of rounds. C2:2** and C2:7** are identical to C2:2 and C2:7, except that the signal about the actions of agents in Country 1 is framed as a standard public signal about θ_1.[24]

3.2.1. *External Validity*

Laboratory experiments are a convenient way to study questions related to information processing, since information sets are typically not observable in naturally occurring data. However, it is important to discuss issues of external validity. To account for experience,

[23]To avoid framing effects, instead of telling subjects in Country 2 the number of people that took action B in Country 1 (as in the model), subjects were given one of the following 3 signals: "0 chose action A, 2 chose action B;" "1 chose action A, 1 chose action B;" or "2 chose action A, 0 chose action B."

[24]In these treatments, the public signal about θ_1 was framed similarly to the private signal about θ_2, that is, as either low, medium, or high.

TABLE VI

EXPERIMENTAL TREATMENTS AND EQUILIBRIUM PREDICTIONS

Treatment	Induced prior	Correlation of states (s)	Precision of y (α)	Location	# subjects	# rounds	Equilibrium actions
C1:1	Optimistic	–	–	NYU	24	30	Roll over for all signals
C1:2	Pessimistic	–	–	NYU	24	30	Roll over for $x = m$ and $x = h$
C2:1	Optimistic	None (1/3)	Uninformative (1/3)	NYU, UCSD	40	30	Roll over for all signals
C2:2	Optimistic	High (3/4)	Uninformative (1/3)	NYU, UCSD	46	30	Roll over for all signals
C2:2*	Optimistic	High (3/4)	Uninformative (1/3)	UCSD	20	50	Roll over for all signals
C2:2**	Optimistic	High (3/4)	Uninformative (1/3)	UCSD	20	30	Roll over for all signals
C2:3	Optimistic	None (1/3)	High (3/4)	NYU, UCSD	48	30	Roll over for all signals
C2:4	Optimistic	High (3/4)	High (3/4)	NYU, UCSD	44	30	Roll over for all signals
C2:5	Optimistic	Perfect (1)	Perfect (1)	NYU, UCSD	48	30	Roll over for all signals
C2:6	Pessimistic	None (1/3)	Uninformative (1/3)	NYU, UCSD	40	30	Roll over for all signals
C2:7	Pessimistic	High (3/4)	Uninformative (1/3)	NYU, UCSD	44	30	Roll over for $x = m$ and $x = h$
C2:7*	Pessimistic	High (3/4)	Uninformative (1/3)	UCSD	24	50	Roll over for $x = m$ and $x = h$
C2:7**	Pessimistic	High (3/4)	Uninformative (1/3)	UCSD	20	30	Roll over for $x = m$ and $x = h$
C2:8	Pessimistic	None (1/3)	High (3/4)	NYU, UCSD	46	30	Roll over for all signals
C2:9	Pessimistic	High (3/4)	High (3/4)	NYU, UCSD	44	30	Roll over for $y = 0$, $y = 1$ & $x = h$
C2:10	Pessimistic	Perfect (1)	Perfect (1)	NYU, UCSD	48	30	Roll over for $y = 0$

subjects in this experiment have an opportunity to learn from the environment and their decisions by playing a large number of repetitions of the game (30 or 50 rounds) with full feedback. Most of the data analysis is based on the last 20 rounds of the experiment (for the 30-round sessions) to allow behavior to stabilize after the learning that takes place in the first 10 rounds. In order to ensure robustness of the results, all sessions from the main set of treatments were run with two different subject pools in different locations, which led to the same qualitative (and similar quantitative) findings. Likewise, a set of robustness checks shows systematic deviations that allow us to conclude that the observed departures from the theory are not due to chance or mistakes.

In terms of portability of the experimental results, a series of papers have compared the performance of undergraduates and professionals in the same experimental setup. In the vast majority of cases, there is no qualitative difference in the types of decisions made by these two groups (see Frechette (2015, 2016) for a survey).[25] This suggests that laboratory experiments can provide a qualitative guide to the type of behavior that might arise in markets. It is important to emphasize the difference between qualitative and quantitative findings (see Kessler and Vesterlund (2015)). For example, systematic biases in the use of information, such as the ones reported in this paper, are qualitative findings that could be common across students and professionals; however, the magnitude of these biases could differ across populations. For this reason, the data analysis will focus on qualitative results by studying systematic deviations from theoretical predictions based on significance and sign, without putting much emphasis on magnitudes of coefficients.[26]

The use of laboratory experiments to address macroeconomic questions is gaining relevance due to the widespread use of explicit microfounded models. As explained by Duffy (2016), we require experimental designs that simplify macroeconomic environments to their core to overcome the complexity of the models. In order to study how subjects process information related to fundamentals and social learning, I simplify the environment as much as possible while keeping the tension between strategic and fundamental uncertainty that characterizes episodes of contagion. Similar to other coordination game experiments, this experiment can be given both a macroeconomic interpretation—-financial contagion—and a microeconomic one, for example, contagion of political revolts. Even with a neutral framing, we observe that the generic forces in the experiment (which will be detailed below) are consistent with what we observe in episodes of contagion, suggesting that the insights that we gain from studying this specific set of microfoundations experimentally can enrich our understanding of episodes of contagion. For example, the contagion to South Korea that followed from the crisis that originated in Thailand in 1997 was characterized by weak fundamental links (consistent with base-rate neglect that led financial actors to ignore the lack of fundamental correlation between the countries) and

[25] In a closely related context, Alevy, Haigh, and List (2007) compared the behavior of students and market professionals from the floor of the Chicago Board of Trade in a game of information cascade formation. Earnings and the rate of cascades do not differ across subject pools, suggesting that insights about social learning behavior can be extended from the student population to professionals. Moreover, they find that professionals are slightly less Bayesian than students, suggesting that the extensive training and experience of traders do not guarantee that their behavior will be consistent with Bayes' rule. This is of particular relevance here, since this paper documents two information processing biases that also imply departures from Bayesian rationality.

[26] The different treatments of the experiment are designed to portray extreme cases where certain signals should or should not be taken into account. For example, when $s = 1/3$ or $\alpha = 1/3$ the social signal should simply not be taken into consideration by subjects in Country 2, because it is uninformative about their payoff-relevant state. These treatments were designed to identify departures in the use of information where magnitudes are not even a question because this signal is simply unrelated to the decision.

anecdotal evidence of overreaction to the observation of what was happening in neighboring countries (consistent with the overreaction to social information).

4. EXPERIMENTAL RESULTS

I briefly describe the behavior of subjects in Country 1 sessions. Then I study the behavior of subjects in Country 2 sessions by testing two main hypotheses about the use of information related to fundamentals and to social behavior and their effects in terms of welfare. Unless otherwise specified, the results will focus on the last 20 rounds played by subjects, to allow for behavior to stabilize.

4.1. *Country 1*

The main question when analyzing Country 1 data is whether subjects behave differently when given an optimistic prior than when given a pessimistic prior. Indeed, there are significant differences in these two treatments and in the expected direction. For each signal observed, there is a significantly lower proportion of decisions to roll over when the prior about the state in Country 1 is pessimistic than when it is optimistic. For each treatment, Table VII shows the percentage of total decisions to roll over for each private signal observed in the last 20 rounds of the experiment (pairwise comparisons across treatments are all statistically different at the 1% level of significance).

This result is in line with the qualitative prediction of more decisions to roll over under optimistic priors. We also observe a higher rate of default for the intermediate state $\theta_1 = M_1$ with the pessimistic prior than with the optimistic prior (61.67% compared to 16.67%), at the 1% level of significance.[27]

Subjects in Country 1 sessions behave qualitatively similar to the existing experimental results on standard global games (see Heinemann, Nagel, and Ockenfels (2004) and Szkup and Trevino (2019)). Across treatments, 85% of subjects use unique monotone strategies, and for each treatment about 50% of total strategies coincide with the equilibrium prediction.

Even if we do not observe all subjects in Country 1 following the equilibrium strategies (rolling over for all signals when facing an optimistic prior; rolling over for medium and high signals, and withdrawing for low signals, when facing a pessimistic prior), we observe a significant difference in behavior in the direction prescribed by the theory, implying a meaningful treatment effect.[28]

TABLE VII

PERCENTAGE OF ROLLOVER DECISIONS, BY SIGNAL AND TREATMENT, C1

Optimistic prior			Pessimistic prior		
Low	Medium	High	Low	Medium	High
31.97%	100%	100%	13.15%	61.83%	92.65%

[27]Note that the relevant state to use in studying rates of default is the intermediate state, $\theta_1 = M_1$, where default occurs when agents miscoordinate their actions.

[28]One could argue that the deviations of the observed behavior of Country 1 subjects from the equilibrium prediction could alter the equilibrium predictions for subjects in Country 2. I calculate the optimal strategies

4.2. *Country 2*

To analyze the behavior of subjects in Country 2 sessions, we focus on how they use the information related to the fundamental and social learning channels of contagion. Recall that subjects observe three pieces of information and, depending on the treatment, all or only a subset of them are relevant for their decision: (1) the prior about θ_1 (optimistic vs. pessimistic), which is related solely to the fundamental channel of contagion, (2) the public signal about the behavior of agents in Country 1 ($y \in \{0, 1, 2\}$), which is related to the social learning channel of contagion, and (3) the private signal about θ_2 ($x_2^i \in \{l_2, m_2, h_2\}$), which, by design, is always informative. I refer to cases where the signals corresponding to the two channels of contagion should determine subjects' behavior as situations where would we expect the channels of contagion to be active.

HYPOTHESIS 1: The fundamental channel of contagion is active whenever the states are correlated ($s > 1/3$). The social learning channel of contagion is active when the states are correlated ($s > 1/3$) and the signal about the behavior in Country 1 is informative ($\alpha > 1/3$).

Table VIII shows the results of five random-effects logit regressions that test Hypothesis 1. For all regressions, the dependent variable is the probability of rolling over. The

TABLE VIII

LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS, BY TREATMENT[a]

	1	2	3	4	5
	$s = 1/3$, $\alpha = 1/3$	$s = 3/4$, $\alpha = 1/3$	$s = 1/3$, $\alpha = 3/4$	$s = 3/4$, $\alpha = 3/4$	$s = 1$, $\alpha = 1$
x_2^i (private signal)	3.035	3.761	3.712	2.932	2.243
	(0.31)	(0.377)	(0.367)	(0.322)	(0.294)
y_{rollover} (social signal)	**0.593**	**0.644**	**1.17**	1.679	1.028
	(0.207)	(0.241)	(0.249)	(0.26)	(0.314)
d_{prior} (prior dummy) (0 opt, 1 pess)	−0.329	**−0.83**	−0.243	**−0.339**	−2.767
	(0.796)	(0.722)	(0.753)	(0.696)	(0.756)
$d_{\text{prior}} * x_2^i$	0.329	**−0.157**	−0.49	**0.193**	0.157
	(0.451)	(0.472)	(0.449)	(0.434)	(0.37)
$d_{\text{prior}} * y_{\text{rollover}}$	**0.556**	0.089	0.406	**−0.271**	0.997
	(0.306)	(0.317)	(0.331)	(0.35)	(0.391)
C	−2.281	−2.147	−3.167	−2.714	−1.367
	(0.562)	(0.533)	(0.575)	(0.514)	(0.598)
N	1600	1800	1858	1756	1836

[a]Clustered (by subject) standard errors in parentheses. Location dummies included in all specifications (see Table XI). Coefficients in bold represent qualitative departures from theoretical predictions.

for agents in Country 2 under the belief that agents in Country 1 behave as in Table VII and call these the "empirical" equilibrium strategies. This changes the equilibrium of only two Country 2 treatments: When $(s, \alpha) = (3/4, 3/4)$, the "empirical" equilibrium strategy is to roll over for $x_2^i = h_2$ and $x_2^i = m_2$ & $y = 0$ and to withdraw otherwise; when $(s, \alpha) = (1, 1)$ agents roll over for $x_2^i = m_2$ & $y = 0$ and $x_2^i = h_2$ & $y = 0$ and withdraw otherwise. The fraction of subjects that use these strategies for the relevant Country 2 treatments is 5% and 12.5%, respectively. In what follows, the analysis of Country 2 takes this into account, if relevant.

independent variables are the private signals x_2^i, the public signal about the number of agents that rolled over in Country 1, y_{rollover}, a dummy variable d_{prior} that takes a value of 0 for an induced optimistic prior and a value of 1 for an induced pessimistic prior; and two interacted terms of this dummy, one with the private signal ($d_{\text{prior}} * x_2^i$) and the other with the public signal ($d_{\text{prior}} * y_{\text{roll}}$) to account for any additional variation of the signals x_2^i and y_{rollover}, respectively, under a different prior. Location dummies are also included in all specifications to account for possible differences in behavior across the two locations where the experiment was run (see Table XI in the Appendix for coefficients of these variables). The five specifications differ in the combination of the parameters (s, α) that define each treatment (see Table VI). In each of these specifications, I pool the data from sessions where an optimistic and a pessimistic prior were induced. I test whether there is a significant difference in behavior under these priors by looking at the coefficient of the dummy d_{prior} and its interacted terms. The coefficients in bold indicate departures, in terms of significance, from the expected results stated in Hypothesis 1. These departures occurred either because coefficients that should be significant are not, or because coefficients that should not be significant are significant.[29] Table XII in the Appendix contains the marginal effects of these regressions.

Table VIII shows that we cannot establish full support for Hypothesis 1 in the data. For both locations, the observed departures from these predictions imply two systematic biases in information processing that are directly related to the two channels of contagion. The coefficients in bold help to identify these biases by illustrating departures, in terms of significance, from theoretical predictions.[30]

One of these biases is related to social learning. In Table VIII, we observe that subjects take into consideration the signal about the behavior of subjects in Country 1 in all the treatments. This signal should have been taken into consideration only in the cases where it is informative ($\alpha > 1/3$) and the states are correlated ($s > 1/3$). We see, for the coefficients in bold, that this signal is a significant determinant of choices and we identify a bias of overreaction to this signal for the cases where it is completely uninformative, either because it is uncorrelated to the true behavior of subjects in Country 1 ($\alpha = 1/3$) or because the states are uncorrelated ($s = 1/3$), in which case it carries no relevant information for subjects in Country 2, irrespective of its precision.[31] For simplicity, in the remainder of the paper I will refer to this overreaction bias as the social imitation bias.[32]

[29]I used departures in terms of significance of coefficients, and not their magnitudes, to focus on qualitative results in order to minimize external validity concerns.

[30]In particular, coefficients in bold relating to y_{rollover} indicate cases where this variable should not have an effect, but instead is found to be significant to, at least, the 5% level. This illustrates the overreaction to the social signal that strengthens the social learning channel. On the other, the coefficients in bold relating to d_{prior} (and its interacted terms) refer to cases where it should be significant, but it is not. This illustrates the base-rate neglect that weakens the fundamental channel.

[31]The reason for the coefficient for y_{rollover} being larger in specification 3 than in specifications 1 and 2 could be that this signal has some accuracy when $\alpha = 3/4$, unlike when $\alpha = 1/3$. A subject might use a heuristic where he partially follows the actions of others if he believes that they knew how to play a similar game (even if uncorrelated to the game he is playing). This cannot be the case when $\alpha = 1/3$, because this signal is not informative about the actions of others. If this is the case, then we "lose" those subjects that try to use such a heuristic when $\alpha = 1/3$, leading to a lower coefficient of y_{rollover}. A similar argument can be used to rule out the interpretation that the overreliance on social information is due to pure imitation, since the social signal does not convey information about the actions of others when $\alpha = 1/3$. Moreover, in all the specifications with an overreaction bias, subjects rely more on their private signal than on the social information, implying that it cannot be pure imitation.

[32]I call this a social imitation bias, and not a social learning bias, because subjects cannot *learn* about their payoff-relevant state from this signal when $s = 1/3$ or $\alpha = 1/3$.

The other bias in information processing that we can identify from Table VIII is related to the fundamental channel in the form of base-rate neglect, since subjects do not take into consideration the prior about θ_1 in situations where the states are correlated ($s > 1/3$). In this case, subjects underweight or ignore the information contained in the prior. We identify the emergence of this bias qualitatively by the lack of statistical significance of the coefficients of the dummy that differentiates the two treatments according to the prior, d_{prior}.[33] Regardless of the social signal, subjects in Country 2 should take the information about the prior into account when the fundamentals are positively correlated ($s > 1/3$). This means that the coefficient of d_{prior} should be statistically different from zero in specifications 2, 4, and 5. Notice that this is the case only for specification 5.[34,35] In specifications 2 and 4, where $s = 3/4$, we see no statistical difference in behavior from the sessions where different priors were induced. This effectively means that even when there is a 75% chance of θ_2 to coincide with θ_1, subjects ignore this information when forming beliefs about θ_2. This is a clear example of base-rate neglect, first introduced by Kahneman and Tversky (1973) and reviewed by Bar-Hillel (1980, 1990), which has been largely documented in individual decision-making tasks.[36] The qualitative results presented in Table VIII remain after controlling for risk aversion (see Table XIII in the Appendix).[37]

The emergence of these biases shows that subjects use public information differently, depending on its source. Both the prior and the social signal are public information, but when this information is presented as a prior, subjects tend to underweight (or completely ignore) it, even if it is informative. However, when presented as a signal based on social behavior, subjects take it into consideration even in cases where it is completely uninformative. Section 4.2.2 contains robustness checks that deepen our understanding of the intrinsic differences between these two types of public information.

With these results in hand, we turn our attention to the second hypothesis, which addresses welfare implications of these biases.

HYPOTHESIS 2: The frequencies of withdrawals and the mean payoffs of subjects in the experiment are consistent with the equilibrium predictions, given the realization of states and signals in the experiment.

[33]One could alternatively identify a base-rate neglect bias by looking at differences in the coefficient of d_{prior}, conditional on the realization of the private signals. Just as in the case of the social imitation bias, I focus on a qualitative identification of biases based on the significance of coefficients, rather than their magnitudes, to minimize issues of external validity. One could posit different hypotheses for why base-rate neglect emerges, such as lack of experience or cognitive load. However, in order to focus on the effect of biases on action choices, we are agnostic about their underlying individual causes.

[34]The fact that base-rate neglect is present when both signals are noisy but highly informative (specification 4), and it is not when both signals are perfectly informative (specification 5) might suggest a higher likelihood of observing this bias as the uncertainty in the environment increases. Since specification 5 implies no base-rate neglect, we can conclude that the emergence of this bias is not a result of experimental procedures, which were identical across treatments.

[35]Notice that in specification 5 subjects overrely on the private signal with respect to the social signal, according to equilibrium. Since we defined the social imitation bias as the case where subjects take into consideration the actions of others when they are irrelevant, this observation is not contradictory of the definition of this bias.

[36]Notice that simple departures from equilibrium (withdrawing when equilibrium suggests rolling over and vice versa) are not necessarily caused by these two biases. Such equilibrium departures can be due to a nonequilibrium use of private information, as in Country 1, where a social imitation bias is not possible and there is no base-rate neglect.

[37]The risk aversion coefficient in Table XIII corresponds to the switching point of the risk measure of Holt and Laury (2002).

For the realization of signals in each treatment, the equilibrium frequency of withdrawals is calculated by looking at the instances where the equilibrium action is to withdraw. In order to understand the welfare consequences of departures from the equilibrium predictions, I construct constrained efficiency and first-best allocation benchmarks. The constrained efficiency benchmark corresponds to a social planner that faces the same informational constraints as the agents and recommends actions consistent with the cooperative solution of the game.[38] In the first-best allocation benchmark, the social planner faces no informational constraints, so the planner's recommendation is for agents to withdraw when the state is low ($\theta_2 = L_2$) and to roll over when the state is either medium or high ($\theta_2 \in \{M_2, H_2\}$).

For each treatment, Panel A of Table IX compares the observed frequency of withdrawals to the equilibrium predictions and to the two efficiency benchmarks. The numbers in parentheses are indices corresponding to the proportion of instances where the observed withdrawals coincide with withdrawals under constrained efficiency and first-best allocations. The higher this index, the smaller the loss in terms of welfare. This information is presented graphically in Figure 2. For the treatments with a pessimistic prior and $(s, \alpha) \in \{(3/4, 3/4), (1, 1)\}$, Table IX shows two additional equilibrium predictions and constrained efficiency benchmarks, which correspond to the "empirical" equilibrium discussed in Footnote 29, that is, the equilibrium prediction for Country 2 that would arise under the belief that agents in Country 1 behave exactly as the Country 1 subjects do in the experiment (see Table VII), which differs from the standard equilibrium prediction for only those two treatments. All numbers in Table IX that refer to the equilibrium, constrained efficiency, and first-best benchmarks correspond to the frequencies of withdrawals or mean payoffs that would arise if we imposed the strategies recommended by these benchmarks to the specific realizations of states and signals in the different treatments of the experiment.

The results in Panel A of Table IX do not support Hypothesis 2. In all of the treatments, the observed frequency of withdrawals is statistically different from the equilibrium predictions (p-values < 0.01 for all but column 7 where p-value < 0.05). Likewise, the observed proportion of withdrawals differs from the constrained efficiency benchmark in all treatments (all p-values < 0.01) and from the first-best allocation in all but two treatments (p-values < 0.01 for columns 2, 6, 7, 8, 10; < 0.05 for columns 5, 9, < 0.1 for column 4, and > 0.1 for columns 1, 3).[39] Notice that in all but two treatments (optimistic prior and $(s, \alpha) \in \{(1/3, 1/3), (1/3, 3/4)\}$), these departures indicate more withdrawals than prescribed by equilibrium. To gain a better understanding of the optimality of the decisions of subjects, we observe the efficiency indices given in parentheses. This is also illustrated by the dark grey filling of the vertical bars in Figure 2, which represents the frequencies of observed withdrawals that coincide with constrained efficiency withdrawals. (The light grey filling represents observed withdrawals that are not prescribed by this benchmark.)

[38]When the prior is optimistic and $s > 1/3$, constrained efficiency coincides with equilibrium predictions (roll over for all signals). When $s = 1/3$, for both an optimistic prior and a pessimistic prior the recommendation is to withdraw when both private signals are low (for all realizations of y), and to roll over otherwise. When the prior is pessimistic and $(s, \alpha) = (3/4, 1/3)$, the recommendation is to withdraw if there is at least one low private signal, and roll over otherwise. When the prior is pessimistic and $(s, \alpha) = (3/4, 3/4)$, the recommendation is to withdraw when $y = 2$ (for all combinations of private signals), when $y = 1$ and at least one private signal is low, and when both private signals are low, and to roll over otherwise. Finally, when the prior is pessimistic and $(s, \alpha) = (1, 1)$, the recommendation is to roll over when $y = 0$, and to withdraw otherwise.

[39]In these two treatments, however, observed withdrawals coincide with those predicted by the first-best allocation less than half of the time.

TABLE IX

COMPARISON OF FREQUENCIES OF WITHDRAWALS AND MEAN PAYOFFS TO EQUILIBRIUM AND EFFICIENCY BENCHMARKS, BY TREATMENT

	Optimistic prior					Pessimistic prior						
	$s=1/3$ $\alpha=1/3$	$s=3/4$ $\alpha=1/3$	$s=1/3$ $\alpha=3/4$	$s=3/4$ $\alpha=3/4$	$s=1$ $\alpha=1$	$s=1/3$ $\alpha=1/3$	$s=3/4$ $\alpha=1/3$	$s=1/3$ $\alpha=3/4$	$s=3/4$ $\alpha=3/4$		$s=1$ $\alpha=1$	
Panel A: Frequency of withdrawals												
Observed (%)	32.5	26.9	31.8	21.3	18.5	32.8	43.8	41.1	41.8		54.1	
Equilibrium (%)	0	0	0	0	0	0	41.1	0	58	58.4†	73.8	81.8†
Cons. eff. (%)	16	0	12.7	0	0	14.3	62.3	16.4	55.7	51.1†	73.8	81.9†
Observed/CE	(0.36)	–	(0.3)	–	–	(0.29)	(0.86)	(0.32)	(0.8)	(0.74)†	(0.88)	(0.96)†
First-best (%)	31.5	19.6	30.2	23.6	15	40	50.5	34.8	46.4		60	
Observed/FB	(0.48)	(0.33)	(0.47)	(0.36)	(0.35)	(0.56)	(0.63)	(0.48)	(0.61)		(0.73)	
Panel B: Mean payoffs (in dollars)												
Observed	10.2	12.57	10.33	12.99	15.15	8.88	7.18	9.09	8.08		6.27	
Equilibrium	13.59	16.09	13.93	15.27	17	12	8.12	12.99	7.43	6.69†	6.37	5.9†
Cons. eff.	13.7	16.09	13.96	15.27	17	12.17	8.31	13.04	8.73	9.27†	6.37	6.28†
Observed/CE	(0.88)	(0.78)	(0.74)	(0.85)	(0.89)	(0.73)	(0.86)	(0.7)	(0.93)	(0.87)†	(0.98)	(0.99)
First-best	14.96	16.87	15.17	16.22	17.6	13.6	11.93	14.44	12.58		10.4	
Observed/FB	(0.68)	(0.75)	(0.68)	(0.8)	(0.86)	(0.65)	(0.6)	(0.63)	(0.64)		(0.6)	

†: predictions for observed behavior in C1 sessions.

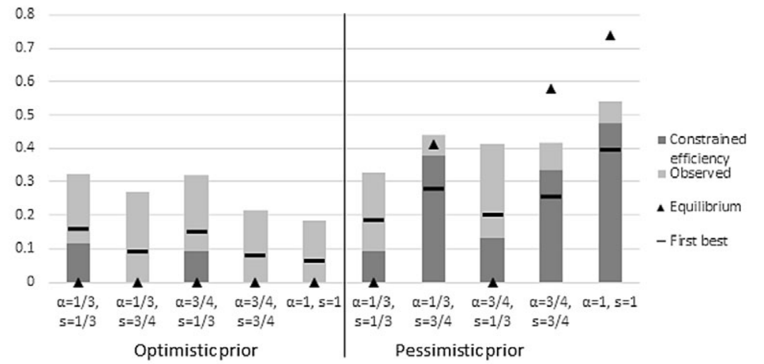

Note: Constrained efficiency (dark grey filling) corresponds to the fraction of observed withdrawals that are predicted by constrained efficiency, not to the total fraction of withdrawals in the constrained efficiency benchmark. The horizontal lines depicting the first best allocation follow the same logic.

FIGURE 2.—Frequencies of withdrawals, by treatment.

On average, there are more discrepancies when the prior is optimistic than when it is pessimistic. The horizontal line in each vertical bar corresponds to the first-best allocation. The discrepancies between the equilibrium frequencies of withdrawals and those corresponding to the first-best allocation illustrate the well-known result regarding the inefficiency of the equilibrium in global games (see Morris and Shin (2003)).

Another indicator of welfare based on performance is mean payoffs. For the observed states and signals in each treatment of the experiment, Panel B in Table IX reports the mean realized payoffs of subjects in the experiment and the mean realized payoffs that would arise under the equilibrium strategies and our two benchmarks. The numbers in parentheses are a measure of relative performance and correspond to the ratio of mean realized payoffs to the payoffs that would arise under the benchmarks. We can interpret these numbers as an index of welfare: The larger the number, the smaller the loss in terms of welfare. Figure 3 presents these results graphically. The loss in terms of welfare with respect to constrained efficiency is illustrated by the dark grey filling of the vertical bars, which corresponds to additional payoffs not realized in the experiment. The horizontal lines correspond to the mean payoffs that would arise under the first-best allocation.

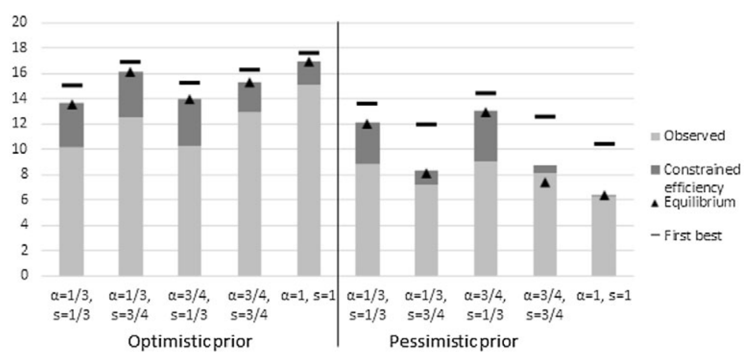

Note: Dark grey filling of bars corresponds to additional payoffs not realized in the experiment that would result under the constrained efficiency benchmark.

FIGURE 3.—Mean payoffs, by treatment.

We can see from Panel B in Table IX that in all but two cases the mean realized payoffs in the experiment are statistically lower than the mean payoffs that would have arisen if subjects had followed the equilibrium strategies (p-values < 0.01, except for the last column where p-value > 0.1). In all but one case, they are lower than the constrained efficiency payoffs (p-values < 0.01, except for the last column where p-value > 0.1), and in all cases they are lower than those corresponding to the first-best allocation (all p-values < 0.01). Subjects realized significant welfare losses by departing from the equilibrium strategies in the treatments where there were more withdrawals than prescribed by equilibrium. However, in one of these treatments (pessimistic prior, $(s, \alpha) = (3/4, 3/4)$) there was a significant gain in terms of welfare from departing from the equilibrium strategy, while in another (pessimistic prior, $(s, \alpha) = (1, 1)$) the realized payoffs of subjects are not statistically different from the equilibrium prediction.[40]

We can interpret these welfare effects as a result of the information processing biases that violate Hypothesis 1. Since there is a strong and systematic overreaction to the social signal, the next subsection analyzes how this overreaction could be related to the specific departures that we observe.

4.2.1. *Social Imitation: The Cost of Following Others*

We now look at the cost, in terms of forgone payoffs, that results when subjects take the action opposite to the one prescribed by equilibrium after they observe a signal that indicates that at least one or exactly two of the subjects in Country 1 took that action. We can identify two types of departure due to the social imitation bias: (1) withdrawing after observing a signal that agents in Country 1 withdrew, when the equilibrium action is to roll over, and (2) rolling over after observing a signal that agents in Country 1 rolled over, when the equilibrium action is to withdraw. Departure (1) can be interpreted as a contagious panic, since subjects react negatively in their own country after observing signals of distress in a foreign country. Departure (2) can be interpreted as a case of contagious confidence, since subjects take actions that imply confidence in their country after observing subjects in a foreign country showing confidence in their own market.

I analyze four cases. The first two correspond to departure (1), and the third and fourth correspond to departure (2). The first case is when subjects observe a signal of at least one subject in Country 1 withdrawing and they withdraw as well, instead of following equilibrium by rolling over. The second case is the same as the first, except that the signal they observe says that both subjects from Country 1 withdrew. The third and fourth cases are analogous to the first two, but for signals about subjects in Country 1 rolling over and as a result they roll over, when the equilibrium action prescribes withdrawing, given the signals that they observe.

[40]We can also analyze whether the equilibrium strategy is optimal, given the distribution of subjects in the sample. We can distinguish five types of subjects: (1) the theoretical equilibrium subjects, (2) the "empirical" equilibrium subjects, who believe that agents in Country 1 behave exactly as the subjects in Country 1, (3) the subjects who take into account the correct information (according to s and α) but do not choose actions according to (1) or (2), (4) the subjects who exhibit the biases that violate Hypothesis 1, and (5) the subjects who behave randomly. There is considerable heterogeneity in the distribution of types within and across treatments. In four treatments, the theoretical equilibrium strategy is still payoff maximizing, given the observed distributions of types, and the payoff losses from equilibrium departures for subjects in categories (3) and (4) were similar. That is, the cost of misusing correct information is not different from the cost of holding incorrect information in the first place.

For each treatment, Table X shows the mean payoffs associated with each of these four cases as well as the mean payoff that would have resulted from following the equilibrium strategy in those instances. Table X also shows the percentage of such instances with respect to the total number of decisions in the last 20 rounds of each treatment, as well as the predictions of the "empirical" equilibrium for the treatments with a pessimistic prior and $(s, \alpha) \in \{(3/4, 3/4), (1, 1)\}$ (see Footnote 29).

We see from Table X that in all treatments where an optimistic prior is induced, and in those where a pessimistic prior is induced and the states are uncorrelated ($s = 1/3$), there is a significant payoff loss from following withdrawals in Country 1 (p-value < 0.01 for all cases). This is consistent with the results of Table IX, which show more instances of withdrawals for those same treatments. The payoff difference in Table X also explains the lower mean realized payoffs for those treatments reported in Table IX. This illustrates the potentially negative effects of the social learning channel of contagion through the emergence of irrational panics. However, when there is a pessimistic prior, the states are correlated ($s \in \{3/4, 1\}$) and the signal about behavior in Country 1 is informative ($\alpha \in \{3/4, 1\}$), we see a significant increase in payoffs resulting from contagious confidence due to the social imitation bias (p-values < 0.01).[41] These two treatments are the only ones that do not exhibit a higher rate of withdrawals than prescribed by equilibrium, as shown in Table IX. This biased increase in confidence suggests that subjects increased their payoffs as a consequence of more successful coordination and fewer withdrawals.[42]

4.2.2. *Robustness of Results*

This section presents new experimental observations to check the robustness of the results presented thus far. The first one asks whether the biases we observe are persistent over time. The second one addresses the framing of the social information.

4.2.2.1. *Longer Sessions.* To better understand the nature of the information processing biases documented here, we explore what happens when subjects play the game for 20 additional rounds (50 in total). The objective of this exercise is to understand whether these biases persist even when subjects have more time to learn how to behave in this environment. For this purpose, I run two additional sessions for the case where $s = 3/4$ and $\alpha = 1/3$, one with an optimistic prior and one with a pessimistic prior. This particular treatment is the natural choice since when $s = 3/4$ and $\alpha = 1/3$ we identify both base-rate neglect (no significance in the coefficient of the prior) and overreaction to the social information (significance of the signal about the actions of others). Table XIV in the Appendix shows the results of a regression analogous to Specification 2 in Table VIII for

[41]The direction of the observed departures for the treatments with an optimistic prior (more withdrawals) is not surprising, since equilibrium predicts that subjects should always roll over. Notice, however, that departures in both directions are a possibility for some treatments with a pessimistic prior, and we in fact see them. Intuitively, observing contagious confidence when the prior is pessimistic (as in the treatments where $(s, \alpha) \in \{(3/4, 3/4), (1, 1)\}$) is a starker result than under an optimistic prior, since subjects are ex ante less likely to believe that rolling over would lead to a preferable outcome when the prior is pessimistic than when it is optimistic.

[42]The fact that social imitation leads to welfare gains with respect to equilibrium illustrates the inefficient nature of the equilibrium in a global game. However, the mean realized payoff resulting from these deviations is still statistically lower than the constrained efficiency benchmark. Thus in these situations, departures from equilibrium lead to welfare gains, but they do not correct all inefficiencies.

TABLE X

PAYOFFS (IN DOLLARS) FROM FOLLOWING OTHERS, BY TREATMENT[a]

	Optimistic prior					Pessimistic prior						
	$s=1/3$ $\alpha=1/3$	$s=3/4$ $\alpha=1/3$	$s=1/3$ $\alpha=3/4$	$s=3/4$ $\alpha=3/4$	$s=1$ $\alpha=1$	$s=1/3$ $\alpha=1/3$	$s=3/4$ $\alpha=1/3$	$s=1/3$ $\alpha=3/4$	$s=3/4$ $\alpha=3/4$		$s=1$ $\alpha=1$	
Follow 1 withdraw	4	4	4	4	4	4	4	4	4	4†	–	–†
Eq. (roll over)	11.02	13.65	10.34	12.91	10.22	9.24	10.31	10.97	5.7	8.57†	–	–†
% total decisions	21.3%	19%	20.2%	14.4%	10.5%	22.4%	7.3%	35%	0.8%	2.4%†	–	–†
Follow 2 withdraw	4	4	4	4	4	4*	4	4	–	4†	–	–†
Eq. (roll over)	12.62	14.32	10.67	12.83	6.25	9.81	7.06	12.32	–	9.3†	–	–†
% total decisions	11.7%	10%	6.8%	6%	1.8%	12.1%	3.9%	18.6%	–	1.7%†	–	–†
Follow 1 roll over	–	–	–	–	–	–	4.31	–	8.74	8.48†	7.09	6.64†
Eq (withdraw)	–	–	–	–	–	–	4	–	4	4†	4	4†
% total decisions	–	–	–	–	–	–	5.8%	–	11.7%	15%†	21.5%	25.7%†
Follow 2 roll over	–	–	–	–	–	–	3.89	-	–	7.59†	–	4.32†
Eq. (withdraw)	–	–	–	–	–	–	4	–	–	4†	–	4†
% total decisions	–	–	–	–	–	–	4.1%	–	–	3.3%†	–	4.2%†

[a] p-values in parenthesis, statistical comparison to equilibrium payoffs.

†: predictions for observed behavior in C1.

three subsets of the data from the longer sessions. One subset discards only the first ten rounds (as in the regressions in Table VIII, to avoid noise from early rounds when subjects get acquainted with the game), another one shows the results of rounds 11-30, which are the same rounds as in Specification 2 of Table VIII, and another one takes into account only the last 20 rounds (31–50) to see how subjects play the game in the additional rounds once they gained more experience.

As we can see, the overreaction to social information persists even when subjects play for 50 rounds. The base-rate neglect in rounds 11–30 is still present (as in Table VIII), however, the 20 additional rounds make the coefficient of the prior significant to the 5% level (rounds 11–50) and if we look only at the last 20 rounds (rounds 31–50, the additional rounds of the robustness check) we can see that the coefficient of the prior is significant to the 1% level. These results suggest that the overreaction that strengthens the social learning channel of contagion is more persistent than the base-rate neglect bias that weakens the fundamental channel, which further highlights the different ways in which subjects process public information, depending on its source.

4.2.2.2. *Use of Public Information.* The experimental results show how public information related to fundamentals and to social behavior can affect the propagation of crises in different ways. Put differently, the results show how public information can be under or overeighted, depending on its source. Public information presented as a prior is neglected in cases where it should be relevant, and public information framed as arising from the behavior of others is taken into consideration even when it is irrelevant. In order to fully characterize the behavioral implications of these biases, we ask whether the framing of the social information affected the results. As shown in the continuous model in Section 2, this signal is theoretically equivalent to a standard public signal about θ_1. For this purpose, I ran additional sessions for Country 2 for the case where $s = 3/4$ and $\alpha = 1/3$, one with an optimistic and one with a pessimistic prior, where instead of observing a public signal about the actions of others, subjects observed a public signal about the realization of the state in Country 1. As mentioned earlier, for these parameters we observe both a base-rate neglect bias and an overreaction bias in the original data. Table XV in the Appendix shows the results of a random-effects logit regression for these additional sessions, analogous to Specification 2 of Table VIII. We can see that when we frame the signal about social behavior as a regular public signal, subjects still overreact to it. This is indicative of a systematic overreaction to public information that is not a prior, regardless of how it is framed.[43] However, the change of frame changes the results with respect to the base-rate neglect. While subjects neglect the prior when this signal is framed as reflecting the behavior of others, they take it into consideration when it is framed as a "regular" public signal (the coefficient for d_{prior} is significant in Table XV and not significant in Specification 2 of Table VIII). A possible interpretation for this difference is that it requires more cognitive effort for subjects to interpret a signal about the behavior of other humans than to interpret a piece of information that comes from a computer. The higher cognitive effort induced by the social frame might lead subjects to spend more cognitive resources to try to understand the reasoning of others that could have led to the signal they observe. As a result, they might be more prone to discard other available information, such as the

[43]This systematic overreaction could be a product of the strategic nature of the environment. Subjects might have used this publicly observed information as an aid to coordination, indicating an important difference between this bias and similar biases studied in individual decision-making environments.

prior, resulting in the base-rate neglect bias. Therefore, the emergence of this bias seems to depend on the nature of the other available signals.[44]

Notice that both of these robustness exercises show strong evidence for the overreaction bias, but not for the base-rate neglect bias. These results emphasize how different types of public information can lead to biases with very different characteristics. Public information presented as a prior can be neglected if social information is present, but this bias is corrected over time, while public information presented as a signal about social behavior is overweighted and this bias is persistent over time. This highlights how social information is intrinsically different from other types of public information and that persistent overreaction to it can have important effects on outcomes and welfare.

5. RELATION TO THE LITERATURE

In this section, I focus on the studies that are closely related methodologically to this paper.[45] Some experimental papers study financial contagion, but none with the use of global games. Brown, Trautmann, and Vlahu (2017), Chakravarty, Fonseca, and Kaplan (2014), and Cipriani, Guarino, Guazzarotti, Tagliati, and Fischer (2018) studied contagion with both fundamental links and social learning.[46] In all of these papers, fundamentals are either perfectly correlated or not correlated at all and the behavior of agents in the first market is perfectly observed and does not constitute a treatment variable. Cipriani et al. (2018) found that contagion occurs in the second market only when fundamentals are perfectly correlated. Their results are consistent with the $(s, \alpha) = (1, 1)$ case in this paper (the only comparable treatment) where we do not observe the emergence of biases. The evidence from Brown, Trautmann, and Vlahu (2017) and Chakravarty, Fonseca, and Kaplan (2014) is inconclusive. While Brown, Trautmann, and Vlahu (2017) found evidence of contagion only when fundamentals are perfectly correlated, Chakravarty, Fonseca, and Kaplan (2014) found that contagion occurs both when they are perfectly correlated and when they are independent. Unlike these papers, I study environments where the correlation of fundamentals can be high but not perfect and where the information observed about the actions of others can be noisy. Moreover, by endowing agents with informative private signals, we are able to ensure that subjects always have information about their payoff-relevant state. These features of the design lead to clear predictions about when each of these channels should be at work and enough variability across treatments to study their interaction. As a result, we are able to identify that subjects overreact to social information and simultaneously neglect purely fundamental information.

The results of this paper are also related to the social learning literature.[47] Weizsacker (2010) performed a meta-study of various experiments on social learning, showing that

[44]Another possible explanation is that it might be easier to combine information that is similar, such as the prior and a signal about the realization of the same variable, than to combine information of different kinds, such as the prior of a variable and the actions of players in a game. If this is the case, then the latter might lead to more mistakes on updating.

[45]The results of this paper are related to many strands of the literature, so it would be unrealistic to attempt to do a full review of all the relevant studies. Even if not discussed in detail, there is an extensive literature that identifies possible channels of contagion based on observations of contagious episodes (see Footnote 1).

[46]Cipriani, Gardenal, and Guarino (2013) found evidence of fundamental-based contagion (abstracting from social learning considerations). Cipriani and Guarino (2008) found evidence of contagion due to social imitation (abstracting from fundamental considerations).

[47]The environment in this paper is different from the one in the classic models of social learning where agents make decisions sequentially and there are no payoff interdependencies across agents. See Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) for theoretical papers, and Anderson and Holt (1997) and the literature that followed for experimental studies.

subjects put too much weight on their private signal in situations where it would be rational to herd. In most social learning studies (and similar to Brown, Trautmann, and Vlahu (2017), Chakravarty, Fonseca, and Kaplan (2014), and Cipriani et al. (2018)) subjects perfectly observe the decisions of their predecessors, unlike in the present study. Celen and Kariv (2005) relaxed this assumption by adding some noise to the social signal and find a positive relationship between the level of uncertainty in the environment and the emergence of behavioral anomalies, which is consistent with our findings. Duffy, Hopkins, and Kornienko (2019) and Goeree and Yariv (2015) study how subjects choose to learn about the state in a social learning experiment. In Duffy, Hopkins, and Kornienko (2019), subjects choose between a private signal about the state and a public signal about the actions of their predecessors, and find no particular bias toward private or social information. Subjects in Goeree and Yariv (2015) choose between an informative private signal about the state and a history of play of predecessors who did not choose a private signal. The latter is uninformative social information, similar to the one in this paper. Goeree and Yariv (2015) showed that one-third of subjects choose the uninformative signal and follow it about 90% of the time and identify this behavior as conformity.[48] These studies, just as the present paper, highlight that social information is different from other types of public information.

The overreaction to unrelated information about the actions of others links this paper to the literature on sunspots, which are public signals unrelated to fundamentals that serve as coordination devices and are mostly studied in macroeconomics. Note, however, that sunspots typically serve as coordination devices for games characterized by multiple equilibria. Global games offer an alternative way to resolve coordination problems, since the introduction of private signals leads to a unique equilibrium. In principle, subjects in this experiment should not have relied on sunspots to coordinate, because the games used in the experiment are global games. However, the results show that subjects are prone to biases in information processing and that the observed departures from Bayesian rationality might have led them to partially base their decisions on the observation of sunspots.[49] It is not clear whether the use of these uninformative signals as partial coordination devices is successful, since the reliance on these signals typically leads to more withdrawals than those predicted by equilibrium, that is, to fewer coordination attempts. However, in two treatments they do serve as successful coordination devices and lead to welfare gains.[50]

This paper is part of a growing literature that studies coordination games with asymmetric information experimentally. Heinemann, Nagel, and Ockenfels (2004), Heinemann, Nagel, and Ockenfels (2009) and Szkup and Trevino (2019) are among those who studied global games in the laboratory. Overall, the experimental evidence supports the type of monotonic strategies prescribed by global games, but finds that subjects deviate from

[48]Suzuki, Jensen, Bossaerts, and O'Dougherty (2016) showed the neural mechanisms behind social contagion in a non-strategic environment. They find that subjects who are more prone to social influence are those who also exhibit activity in brain regions associated with the theory of mind (who understand how to play games of strategy). These findings suggest that the subjects exhibiting social imitation might not necessarily be the "irrational" types.

[49]We cannot say that subjects in this experiment rely *fully* on sunspots because the private signal is a strong driver of behavior in all treatments.

[50]Duffy and Fisher (2005) showed experimentally that sunspots can enhance coordination and that the framing of the variable used as a sunspot affects its success as a coordination device. Fehr, Heinemann, and Llorente-Saguer (2019) found that the presence of private signals reduces the power of public signals as sunspots, because subjects cannot ignore their private signal, which is consistent with some of our results.

equilibrium toward more efficient strategies and that the extent of this departure toward efficiency depends on the amount of uncertainty in the environment (Szkup and Trevino (2019)). Figure 2 shows that this is the case in the treatments where the observed behavior of subjects is more efficient than the behavior prescribed by equilibrium. In the related family of beauty contest games with private and public information, Baeriswyl and Cornand (2016) and Cornand and Heinemann (2014) showed that subjects underweight the role of the public signal as a coordination device and use a model with limited depth of reasoning to explain their results. The present paper involves two types of public signals, the prior and the social signal, and thus expands the results of these papers by showing that subjects underweight the former and overweight the latter, suggesting that the type of public signal might determine how it is treated. The robustness check related to framing of the social signal, presented in Section 4.2.2, strengthens this point.[51]

6. DISCUSSION AND CONCLUDING REMARKS

This paper presents the results of an experiment designed to understand how information about fundamental links and social behavior can trigger financial contagion using a global games model as the theoretical benchmark. While the theory provides very clear predictions as to when each of these channels should lead to contagion, the experimental results show systematic departures in the form of two information processing biases that directly affect the mechanism of each channel. Base-rate neglect leads subjects to underweight fundamental correlation, and overreaction to the signal about the behavior of others strengthens the social learning channel. This persistent overreaction leads the social learning channel to operate even in the absence of fundamental correlation, which is a stark departure from the theoretical (rational) benchmark. These results highlight the necessity to study financial crises through a behavioral lens and contribute to the growing literature that is bridging the divide between behavioral economics and macroeconomic models.[52]

One takeaway of this analysis is that not all types of public information are treated equally by agents. The emergence of these biases shows that the source of public information determines how people process it. In Bayesian models, public information reflecting social behavior is equivalent to a public signal coming from a standard data-generating process. However, the results of this experiment show that people overreact to public information about social behavior in a persistent manner. Moreover, people tend to ignore prior information when social information is also present, but not in the presence of a standard public signal. In the context of financial contagion, the presence and overreaction to social information can crucially affect how crises propagate. Taking Bayesian models at face value would ignore this behavioral response to social information, which amplifies the effects of contagion and has significant welfare considerations.

[51]The strong effect of the social imitation bias can also affect the way we think about other mechanisms for contagion, such as adverse selection. Morris and Shin (2012) suggested that even small amounts of adverse selection can unravel market confidence and result in the breakdown of trade in an asset market. The results of Section 4.2.1 can be interpreted as suggesting that, in the presence of adverse selection, the social imitation bias could exacerbate the failure of market confidence that results from adverse selection (if the bias leads to contagious panics), or it could potentially mitigate it (if the bias leads to contagious confidence).

[52]In a similar spirit, Bordalo, Gennaioli, and Shleifer (2018) and Gennaioli, Shleifer, and Vishny (2012) studied the effects of behavioral biases in credit cycles and financial fragility, respectively. See also Barberis (2015) for psychology-based mechanisms as possible forces behind the financial crises of 2007–2008.

In terms of policy implications, the biased response to different types of public signals highlights an intrinsic difference between public information sources that inform economic agents. People might not respond in the same way to government announcements, which can be thought of as standard public signals, and market observations, which are public signals generated by social behavior, even if their informational content is the same. This suggests that these sources of information might in fact not be substitutes, which is an important consideration for public announcements by policy makers.

The departures from Bayesian rationality implied by these biases affect macroeconomic mechanisms leading to significant welfare effects that are not accounted for in standard models. These findings contribute to our understanding of how crises and runs might originate and propagate, and they contrast other mechanisms proposed in the macroeconomics literature that explicitly invoke assumptions on preferences or distributions to explain observed behavior that standard models cannot predict. Model misspecification and ambiguity aversion, for example, study situations where agents are uncertain about their environment and "doubt" the specification of the data-generating process. Hansen and Sargent (2001, 2011) proposed optimal policies to account for the effects of these preferences on outcomes. Alternative assumptions on the data generating process, such as fat-tailed distributions, can also explain observations that are not captured by standard models (see, e.g., Farhi and Gabaix (2016) or Gourio (2012)). Similar to these models, the biases identified in this paper lead to different behavior than that predicted by standard models, and thus to new policy considerations. Unlike these models, the mechanism identified in this paper is based on biased information processing which is rooted in departures from Bayesian rationality that are well documented in the behavioral literature, and not on assumptions.

As with many papers that identify anomalies in decision making, the next step is to think of possible ways to extend the standard model to rationalize the evidence presented here. The development of a satisfactory behavioral model should have strong foundations based on exhaustive empirical evidence to truly advance our understanding of these biases in strategic settings. Developing such empirical and theoretical exercises constitutes an exciting avenue for future research.

One can also think of alternative models that are not based on behavioral biases that could explain some of the findings and potentially allow for the two channels of contagion to operate independently. An example of an alternative equilibrium model is one where rational players face an opponent that is a "follower/panic" type with probability p. This type of player follows the action of his predecessors when they coordinate on the same action, or panics and withdraws when they miscoordinate, regardless of his private signal and of the informativeness of the social signal. As a result, a rational player might also take irrelevant social information into consideration in an effort to coordinate with her opponent if she believes that there is a high enough chance that her opponent is a "follower/panic" type. This type of model could rationalize the experimental results documented in this paper when $p \in [0.28, 0.38]$, depending on the treatment.

A popular nonequilibrium model is level-k (see Crawford, Costa-Gomes, and Iriberri (2013) for a survey). In this case, players can have limited depth of strategic sophistication and/or of signal extraction ability. When players are limited only in their depth of strategic sophistication (as in standard level-k models), this approach could predict the larger frequencies of withdrawals with respect to equilibrium that we observe both in Country 1 and in some treatments of Country 2. However, this is due to anchoring beliefs to level-0 players who uniformly randomize their actions, and not to a biased use of information, which is the key mechanism in this paper.[53] This level-k model would not capture the systematic overreaction to social information, and it would even predict an underreaction for treatments where this information is actually relevant. Alternatively, we can think of a nonstandard level-k model where players are limited in their ability to update information. A level-0 player could naively update beliefs by giving equal weight to all signals, regardless of their precision, which could, in principle, lead higher-level players to take into consideration the social signal (even if irrelevant) in an effort to coordinate with their opponents. This could potentially help explain the overreaction bias, but not the base-rate neglect bias, because such a naive updating rule would imply that level-0 players take all signals into consideration, including the prior.[54,55]

An interesting consideration that arises when thinking about alternative models is to depart from the unique equilibrium prediction that global games typically make. As mentioned in Section 2.2, all the parameters in the experiment were chosen to ensure a unique equilibrium in order to derive sharper theoretical predictions that would help characterize behavioral departures. However, an interesting extension could be to study what happens if we use parameters that lead to multiple equilibria. Departures from uniqueness in a global game model have complex belief dynamics because of the breakdown of common knowledge that characterizes these games, which should be taken into account carefully if one were to study equilibrium selection in global games. Morris and Yildiz (2019) is the first paper to explicitly study equilibrium selection in global games with multiple equilibria, leading to a "second generation" class of global games models. This mechanism is equally fascinating and complex, and using this framework for an experimental investigation is a very interesting endeavor for future research.[56]

[53]Given level-0 behavior, level-1 players believe that there is only a $\frac{1}{2}$ probability of their opponent rolling over—as opposed to, for example, the probability of 1 that equilibrium predicts when the prior is optimistic. This makes them less likely to roll over, and, as a consequence, level-2 players believe that their level-1 opponent rolls over with probability less than 1, and so on, leading to more withdrawals than predicted by equilibrium.

[54]This model would relax a different assumption behind Nash equilibrium than typical level-k models. Because of the complex structure of the game, how to come up with natural assumptions about strategic level-0 players who are naive updaters in Country 2 is not obvious.

[55]We could also think of a model where we relax both assumptions simultaneously, that is, where a level-(j, k) player forms beliefs about the state according to his level of signal-extraction sophistication, j, and about the action of his opponent according to his level of strategic sophistication, k. Such a model could be very appealing but also nontrivial. Certain assumptions (such as rooting all behavior in a level-$(0, 0)$ player) could result in predictions similar to those of its constituent parts where, for example, information processing biases would not play a role in driving the results in the way that they do in the data. Moreover, assumptions about the relationship between levels j and k should have empirical foundations. The development of such a model is beyond the scope of this paper.

[56]In terms of bringing the experiment closer to the macroeconomic environment behind the model, notice that agents that make decisions using standard debt contracts need not do so simultaneously and these contracts typically have rolling over as a default option. The experimental implementation of this paper abstracts from these specificities in order to minimize possible sources of confounds when analyzing how subjects process information related to fundamentals and social learning. Sequential decisions could inject an additional layer of complexity by leading to different belief formation processes, depending on the order of play. Having a default action could lead to a status quo bias, which would imply a larger fraction of rollover decisions. As an extension, it would be interesting to study whether overreaction to social information is strong enough to overcome potential status quo biases.

APPENDIX

TABLE XI

LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS, BY TREATMENT[a]

	1	2	3	4	5
	$s = 1/3$, $\alpha = 1/3$	$s = 3/4$, $\alpha = 1/3$	$s = 1/3$, $\alpha = 3/4$	$s = 3/4$, $\alpha = 3/4$	$s = 1$, $\alpha = 1$
x_2^i	3.035	3.761	3.712	2.932	2.243
	(0.31)	(0.377)	(0.367)	(0.322)	(0.294)
y_{rollover}	**0.593**	**0.644**	**1.17**	1.679	1.028
	(0.207)	(0.241)	(0.249)	(0.26)	(0.314)
d_{prior}	−0.329	**−0.83**	−0.243	**−0.339**	−2.767
(0 opt, 1 pess)	(0.796)	(0.722)	(0.753)	(0.696)	(0.756)
$d_{\text{prior}} * x_2^i$	0.329	**−0.157**	−0.49	**0.193**	0.157
	(0.451)	(0.472)	(0.449)	(0.434)	(0.37)
$d_{\text{prior}} * y_{\text{roll}}$	**0.556**	0.089	0.406	**−0.271**	0.997
	(0.306)	(0.317)	(0.331)	(0.35)	(0.391)
location	0.701	1.51	0.183	0.862	−2.809
(0 NYU, 1 UCSD)	(0.79)	(0.705)	(0.774)	(0.7)	(0.909)
$x_2^i * location$	0.242	−1.989	−0.665	−0.873	0.324
	(0.445)	(0.417)	(0.449)	(0.398)	(0.392)
$y_{\text{rollover}} * location$	−0.22	−0.417	−0.164	−0.391	1.187
	(0.312)	(0.295)	(0.323)	(0.344)	(0.464)
$d_{\text{prior}} * location$	0.738	−0.386	−0.532	−0.197	3.42
(0 opt, 1 pess)	(1.11)	(0.982)	(1.068)	(0.927)	(1.1)
$d_{\text{prior}} * x_2^i * location$	−1.286	0.658	0.262	−0.34	−0.552
	(0.621)	(0.563)	(0.595)	(0.53)	(0.497)
$d_{\text{prior}} * y_{\text{roll}} * location$	−0.657	−0.012	0.156	−0.054	−1.719
	(0.426)	(0.396)	(0.456)	(0.453)	(0.559)
C	−2.281	−2.147	−3.167	−2.714	−1.367
	(0.562)	(0.533)	(0.575)	(0.514)	(0.598)
N	1600	1800	1858	1756	1836

[a]Clustered (by subject) standard errors in parentheses. Coefficients in bold represent qualitative departures from theoretical predictions.

TABLE XII

MARGINAL EFFECTS OF LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS, BY TREATMENT[a]

	1	2	3	4	5
	$s=1/3$, $\alpha=1/3$	$s=3/4$, $\alpha=1/3$	$s=1/3$, $\alpha=3/4$	$s=3/4$, $\alpha=3/4$	$s=1$, $\alpha=1$
x_2^i	0.398	0.546	0.572	0.372	0.384
	(0.073)	(0.078)	(0.075)	(0.049)	(0.051)
y_{rollover}	**0.078**	**0.093**	**0.18**	0.213	0.176
	(0.03)	(0.036)	(0.041)	(0.037)	(0.054)
d_{prior}	−0.043	**−0.122**	−0.038	**−0.043**	−0.451
(0 opt, 1 pess)	(0.105)	(0.109)	(0.117)	(0.089)	(0.12)
$d_{\text{prior}} * x_2^i$	0.043	**−0.023**	−0.075	**0.024**	0.027
	(0.059)	(0.068)	(0.069)	(0.056)	(0.064)
$d_{\text{prior}} * y_{\text{roll}}$	**0.073**	0.013	0.063	**−0.034**	0.171
	(0.041)	(0.046)	(0.052)	(0.044)	(0.07)
location	0.091	0.215	0.028	0.11	−0.458
(0 NYU, 1 UCSD)	(0.104)	(0.101)	(0.118)	(0.091)	(0.137)
$x_2^i * location$	0.032	−0.289	−0.102	−0.111	0.055
	(0.058)	(0.064)	(0.069)	(0.05)	(0.068)
$y_{\text{roll}} * location$	−0.029	−0.06	−0.025	−0.05	0.203
	(0.041)	(0.043)	(0.05)	(0.044)	(0.08)
$d_{\text{prior}} * location$	0.084	−0.06	−0.09	−0.026	0.386
(0 opt, 1 pess)	(0.111)	(0.163)	(0.195)	(0.125)	(0.092)
$d_{\text{prior}} * x_i * location$	−0.169	0.096	0.04	−0.043	−0.095
	(0.083)	(0.081)	(0.091)	(0.068)	(0.086)
$d_{\text{prior}} * y_{\text{roll}} * location$	−0.086	−0.002	0.024	−0.007	−0.295
	(0.057)	(0.057)	(0.07)	(0.058)	(0.098)

[a]Clustered (by subject) standard errors in parentheses. Coefficients in bold represent qualitative departures from theoretical predictions.

TABLE XIII

LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS CONTROLLING FOR RISK AVERSION, BY TREATMENT[a]

	6	7	8	9	10
	$s = 1/3$, $\alpha = 1/3$	$s = 3/4$, $\alpha = 1/3$	$s = 1/3$, $\alpha = 3/4$	$s = 3/4$, $\alpha = 3/4$	$s = 1$, $\alpha = 1$
x_2^i	3.641	3.957	3.987	3.25	2.302
	(0.413)	(0.411)	(0.412)	(0.369)	(0.304)
y_{roll}	**0.63**	**0.681**	**1.25**	1.875	1.048
	(0.246)	(0.254)	(0.263)	(0.292)	(0.32)
d_{prior}	0.025	**−0.922**	−0.155	**−0.24**	−2.953
(0 opt, 1 pess)	(1.029)	(0.688)	(0.939)	(0.875)	(0.819)
$d_{\text{prior}} * x_i$	−0.016	**−0.266**	−0.58	**0.448**	0.215
	(0.537)	(0.487)	(0.497)	(0.527)	(0.386)
$d_{\text{prior}} * y_{\text{roll}}$	**0.598**	0.055	0.418	**−0.263**	1.08
	(0.341)	(0.326)	(0.352)	(0.399)	(0.404)
Risk aversion	0.196	−0.829	−0.036	−0.878	0.031
(Holt–Laury index)	(0.258)	(0.171)	(0.294)	(0.287)	(0.15)
C	−3.893	2.254	−3.205	1.579	−1.524
	(0.665)	(1.008)	(1.464)	(1.578)	(0.971)

[a] Clustered (by subject) standard errors in parentheses. Coefficients in bold represent qualitative departures from theoretical predictions.

TABLE XIV

LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS IN EXPERIMENT WITH 50 ROUNDS[a]

	$s = 3/4, \alpha = 1/3$		
Rounds:	11–30	11–50	31–50
x_2^i	3.025	2.885	2.771
	(0.284)	(0.2)	(0.303)
y_{rollover}	**0.824**	**0.653**	**0.509**
	(0.212)	(0.149)	(0.23)
d_{prior}	**−0.036**	−0.684	−1.652
(0 opt, 1 pess)	(0.463)	(0.321)	(0.493)
$d_{\text{prior}} * x_i$	−0.078	−0.018	0.03
	(0.365)	(0.254)	(0.372)
$d_{\text{prior}} * y_{\text{roll}}$	0.191	0.165	0.189
	(0.284)	(0.198)	(0.294)
C	−2.688	−2.741	−1.769
	(0.42)	(0.35)	(0.431)
N	880	1760	880

[a] Clustered (by subject) standard errors in parentheses. Coefficients in bold represent qualitative departures from theoretical predictions.

TABLE XV

LOGIT ESTIMATES OF INFORMATION TAKEN INTO ACCOUNT FOR INDIVIDUAL ACTIONS IN EXPERIMENT WITH DIFFERENT FRAME[a]

	$s = 3/4$, $\alpha = 1/3$
x_2^i	3.014
	(0.334)
y_{rollover}	**1.067**
	(0.256)
d_{prior}	−2.98
(0 opt, 1 pess)	(0.873)
$d_{\text{prior}} * x_i$	−0.01
	(0.464)
$d_{\text{prior}} * y_{\text{roll}}$	**0.801**
	(0.385)
C	−1.543
	(0.92)
N	800

[a] Clustered (by subject) standard errors in parentheses. Coefficients in bold represent qualitative departures from theoretical predictions.

REFERENCES

AHNERT, T., AND C. BERTSCH (2018): "A Wake-Up Call: Information Contagion and Speculative Currency Attacks," Report. [299]

ALEVY, J., M. HAIGH, AND J. LIST (2007): "Information Cascades: Evidence From a Field Experiment With Financial Market Professionals," *Journal of Finance*, 62, 151–180. [313]

ANDERSON, L., AND C. HOLT (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87 (5), 847–862. [325]

ANGELETOS, G. M., AND I. WERNING (2006): "Crises and Prices: Information Aggregation, Multiplicity, and Volatility," *American Economic Review*, 96 (5), 1720–1736. [299]

ANGELETOS, G. M., C. HELLWIG, AND A. PAVAN (2006): "Signaling in a Global Game: Coordination and Policy Traps," *Journal of Political Economy*, 114 (3), 452–484. [299]

——— (2007): "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks," *Econometrica*, 75 (3), 711–756. [299]

BAERISWYL, R., AND C. CORNAND (2016): "The Predominant Role of Signals Precision in Experimental Beauty Contest," *The B.E. Journal of Theoretical Economics*, 16 (1), 267–301. [327]

BANERJEE, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107 (3), 797–817. [325]

BAR-HILLEL, M. (1980): "The Base Rate Fallacy in Probability Judgements," *Acta Psychologica*, 44, 211–233. [317]

——— (1990): "Back to Base Rates," in *Insights in Decision Making: A Tribute to Hillel J. Einhorn*. Chicago, IL: University of Chicago Press, 200–216. [317]

BARBERIS, N. (2015): "Psychology and the Financial Crisis of 2007–2008," in *Financial Innovation: Too Much or Too Little?* ed. by M. Haliassos. Cambridge, MA: MIT Press. [327]

BASU, R. (2002): "Financial Contagion and Investor 'Learning': An Empirical Investigation," IMF Working Paper 02/218, International Monetary Fund. [298]

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade," *Journal of Political Economy*, 100 (5), 992–1026. [325]

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2018): "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, 73 (1), 199–227. [327]

BROWN, M., S. TRAUTMANN, AND R. VLAHU (2017): "Understanding Bank Run Contagion," *Management Science*, 63 (7), 2272–2282. [325,326]

CALVO, G., AND E. MENDOZA (2000): "Rational Contagion and the Globalization of Securities Markets," *Journal of International Economics*, 51 (1), 79–113. [297]

CARAMAZZA, F., L. RICCI, AND R. SALGADO (2004): "International Financial Contagion in Currency Crises," *Journal of International Money and Finance*, 23 (1), 51–70. [298]

CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61 (5), 989–1018. [298]

CELEN, B., AND S. KARIV (2005): "An Experimental Test of Observational Learning Under Imperfect Information," *Economic Theory*, 26 (3), 677–699. [326]

CHAKRAVARTY, S., M. A. FONSECA, AND T. R. KAPLAN (2014): "An Experiment on the Causes of Bank Run Contagions," *European Economic Review*, 72, 39–51. [325,326]

CIPRIANI, M., AND A. GUARINO (2008): "Herd Behavior and Contagion in Financial Markets," *The B.E. Journal of Theoretical Economics*, 8 (1). [325]

CIPRIANI, M., G. GARDENAL, AND A. GUARINO (2013): "Financial Contagion in the Laboratory: The Cross-Market Rebalancing Channel," *Journal of Banking and Finance*, 37 (11), 4310–4326. [325]

CIPRIANI, M., A. GUARINO, G. GUAZZAROTTI, F. TAGLIATI, AND S. FISCHER (2018): "Informational Contagion in the Laboratory," *Review of Finance*, 22 (3), 877–904. [325,326]

CLAESSENS, S., AND K. FORBES (Eds.) (2001): *International Financial Contagion*. Boston, MA: Kluwer Academic. [297]

CORNAND, C., AND F. HEINEMANN (2014): "Measuring Agents' Overreaction to Public Information in Games With Strategic Complementarities," *Experimental Economics*, 17 (1), 61–77. [327]

CRAWFORD, V., M. COSTA-GOMES, AND N. IRIBERRI (2013): "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 51, 5–62. [328]

DASGUPTA, A. (2004): "Financial Contagion Through Capital Connections: A Model of the Origin and Spread of Bank Panics," *Journal of the European Economic Association*, 2 (6), 1049–1084. [299]

——— (2007): "Coordination and Delay in Global Games," *Journal of Economic Theory*, 134 (1), 195–225. [304]

DUFFY, J. (2016): "Macroeconomics: A Survey of Laboratory Research," in *The Handbook of Experimental Economics*, Vol. 2, ed. by J. Kagel and A. Roth. Princeton, NJ: Princeton University Press. [313]

DUFFY, J., AND E. FISHER (2005): "Sunspots in the Laboratory," *American Economic Review*, 95, 510–529. [326]

DUFFY, J., E. HOPKINS, AND T. KORNIENKO (2019): "Lone Wolf or Herd Animal? An Experiment on Choice of Information and Social Learning," Report. [326]

EDMOND, C. (2013): "Information Manipulation, Coordination and Regime Change," *Review of Economic Studies*, 80, 1422–1458. [299]

FARHI, E., AND X. GABAIX (2016): "Rare Disasters and Exchange Rates," *Quarterly Journal of Economics*, 131 (1), 1–52. [328]

FEHR, D., F. HEINEMANN, AND A. LLORENTE-SAGUER (2019): "The Power of Sunspots: An Experimental Analysis," *Journal of Monetary Economics* 103, 123–136. [326]

FISCHBACHER, U. (2007): "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 10 (2), 171–178. [310]

FRECHETTE, G. (2015): "Laboratory Experiments: Professionals versus Students," in *Handbook of Experimental Economic Methodology*, ed. by G. Frechette and A. Schotter. Oxford: Oxford University Press. [313]

——— (2016): "Experimental Economics Across Subject Populations," in *The Handbook of Experimental Economics*, Vol. 2, ed. by J. Kagel and A. Roth. Princeton, NJ: Princeton University Press. [313]

GENNAIOLI, N., A. SHLEIFER, AND R. VISHNY (2012): "Neglected Risks, Financial Innovation and Financial Fragility," *Journal of Financial Economics*, 104 (3), 452–468. [327]

GOEREE, J., AND L. YARIV (2015): "Conformity in the Lab," *Journal of the Economic Science Association*, 1 (1), 15–28. [326]

GOLDSTEIN, I. (2012): "Empirical Literature on Financial Crises: Fundamentals vs. Panic," in *The Evidence and Impact of Financial Globalization*, ed. by G. Caprio. Amsterdam: Elsevier. [298]

GOLDSTEIN, I., AND A. PAUZNER (2004): "Contagion of Self-Fulfilling Financial Crises Due to Diversification of Investment Portfolios," *Journal of Economic Theory*, 119 (1), 151–183. [299]

GOURIO, F. (2012): "Disaster Risk and Business Cycles," *American Economic Review*, 102 (6), 2734–2766. [328]

HANSEN, L. P., AND T. J. SARGENT (2001): "Acknowledging Misspecification in Macroeconomic Theory," *Review of Economic Dynamics*, 4 (3), 519–535. [328]

——— (2011): "Wanting Robustness in Macroeconomics," in *Handbook of Monetary Economics*, ed. by B. Friedman and M. Woodford. Elsevier Science. [328]
HEINEMANN, F., R. NAGEL, AND P. OCKENFELS (2004): "The Theory of Global Games on Test: Experimental Analysis of Coordination Games With Public and Private Information," *Econometrica*, 72 (5), 1583–1599. [306,314,326]
——— (2009): "Measuring Strategic Uncertainty in Coordination Games," *Review of Economic Studies*, 76 (1), 181–221. [326]
HELLWIG, C. (2002): "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games," *Journal of Economic Theory*, 107 (2), 191–222. [304]
HELLWIG, C., A. MUKHERJI, AND A. TSYVINSKI (2006): "Self-Fulfilling Currecy Crisis: The Role of Interest Rates," *American Economic Review*, 96 (5), 1769–1787. [299]
HOLT, C. A., AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92 (5), 1644–1655. [317]
KAHNEMAN, D., AND A. TVERSKY (1973): "On the Psychology of Prediction," *Psychological Review*, 80, 237–251. [317]
KAMINSKY, G., AND S. SCHMUKLER (1999): "What Triggers Market Jitters? A Chronology of the Asian Crisis," *Journal of International Money and Finance*, 18 (4), 537–560. [298]
KAMINSKY, G. L., R. LYONS, AND S. L. SCHMUKLER (2004): "Managers, Investors, and Crises: Mutual Fund Strategies in Emerging Markets," *Journal of International Economics*, 64, 113–134. [298]
KAMINSKY, G. L., C. M. REINHART, AND C. A. VEGH (2003): "The Unholy Trinity of Financial Contagion," *Journal of Economic Perspectives*, 17 (4), 51–74. [297]
KESSLER, J., AND L. VESTERLUND (2015): "The External Validity of Laboratory Experiments: The Misleading Emphasis on Quantitative Effects," in *Handbook of Experimental Economic Methodology*, ed. by G. Frechette and A. Schotter. Oxford: Oxford University Press. [313]
KODRES, L., AND M. PRITSKER (2002): "A Rational Expectations Model of Financial Contagion," *Journal of Finance*, 57 (2), 769–800. [297]
MANZ, M. (2010): "Information-Based Contagion and the Implications for Financial Fragility," *European Economic Review*, 54 (7), 900–910. [299]
MORRIS, S., AND H. S. SHIN (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88 (3), 587–597. [299]
——— (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, ed. by M. Dewatripont, L. Hansen, and S. Turnovsky. Cambridge: Cambridge University Press. [299,304,320]
——— (2004): "Coordination Risk and the Price of Debt," *European Economic Review*, 48 (1), 133–153. [299, 302,304,305]
——— (2012): "Contagious Adverse Selection," *American Economic Journal: Macroeconomics*, 4 (1), 1–21. [327]
MORRIS, S., AND M. YILDIZ (2019): "Crises: Equilibrium Shifts and Large Shocks," *American Economic Review*, 109 (8), 2823–2854. [329]
OH, F. D. (2013): "Contagion of a Liquidity Crisis Between Two Firms," *Journal of Financial Economics*, 107 (2), 386–400. [299]
SCHMUKLER, S., P. ZOIDO-LOBATON, AND M. HALAC (2006): "Financial Globalization, Crises, and Contagion," in *International Macroeconomics: Recent Developments*, ed. by A. Morales. Hauppauge, NY: Nova Science Publishers. [297]
SUZUKI, S., E. JENSEN, P. BOSSAERTS, AND J. O'DOUGHERTY (2016): "Behavioral Contagion During Learning About Another Agent's Risk-Preferences Acts on the Neural Representation of Decision-Risk," *Proceedings of the National Academy of Sciences*, 113 (14), 3755–3760. [326]
SZKUP, M. (2015): "Debt Overhang and Liquidity Risk," Report. [302]
SZKUP, M., AND I. TREVINO (2019): "Strategic Uncertainty and Information Structures in Coordination Games," Report. [306,314,326,327]
TREVINO, I. (2020): "Supplement to 'Informational Channels of Financial Contagion'," *Econometrica Supplemental Material*, 88, https://doi.org/10.3982/ECTA15604. [310]
WEIZSACKER, G. (2010): "Do We Follow Others When We Should? A Simple Test of Rational Expectations," *American Economic Review*, 100 (5), 2340–2360. [325]

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 18 August, 2017; final version accepted 25 August, 2019; available online 18 September, 2019.