

Sand in the Wheels: A Dynamic Global-Game Approach*

LAURENT MATHEVET[†]

JAKUB STEINER[‡]

University of Texas at Austin

Northwestern University

December 5, 2011

Abstract

We study the impact of frictions on the prevalence of systemic crises. Agents privately learn about a fixed payoff parameter, and repeatedly adjust their investments while facing transaction costs in a dynamic global game. The model has a rich structure of externalities: payoffs may depend on the volume of aggregate investment, on the concentration of investment, or on its volatility. We examine how small frictions, including those similar to the Tobin tax, affect the equilibrium. We identify conditions under which frictions discourage harmful behavior without compromising investment volume. The analysis is driven by a robust invariance result: the volume of aggregate investment (measured in a pivotal contingency) is invariant to a large family of frictions.

*We have benefited from comments by Sylvain Chassang, Federico Echenique, Eugen Kováč, Filip Matějka, David Myatt, Alessandro Pavan, József Sákovics, Ennio Stacchetti, Colin Stewart, Jonathan Weinstein, and from audiences at NYU, HECER in Helsinki, the University of Texas at Austin, at the Kansas Theory Workshop, at the Midwest theory conference in Vanderbilt, and at the conferences “Information and Coordination: Theory and Applications”, and EEA-ESEM in OSLO. Rossella Argenziano and Siyang Xiong provided excellent comments as discussants.

[†]lmath@mail.utexas.edu

[‡]j-steiner@kellogg.northwestern.edu

1 Introduction

The recent economic crisis has led to a renewed interest in Tobin’s proposal to “throw some sand in the wheels” of the economy.¹ One voice in the current public debate points to the stabilizing role of small frictions. The advocates of the transaction tax, such as France or Germany, believe that it would calm the self-fulfilling financial turmoils experienced in recent years. The opponents, such as the United Kingdom or the United States, worry that the tax may reduce financial liquidity or shift the investments outside of the taxing jurisdictions.

We examine the stabilizing role of frictions in a coordination game where both high and low economic activity can become a self-fulfilling prophecy. In our model, the economic turmoils arise as rational reactions to evolving information: investors try to outguess the coordination outcome of the whole economy, and perpetually adjust their investment positions. The frictions influence properties of the adjustment patterns, such as their volatility, thereby impacting real economic performance. In equilibrium, frictions can have a large effect on the prevalence of systemic crises.

We model the coordination of the economy in a dynamic global game. A continuum of agents gradually and privately learn about the fixed state of the economy, and repeatedly adjust their investment positions. An agent’s payoff depends on her investment path and on the outcome of the economy, which either succeeds or fails. The outcome is a function of various statistics of investors’ behavior, including the terminal volume of investment, the volatility of investment, the exit rate, or the dispersion of investment across agents. Allowing for general transaction costs, we examine how frictions impact these statistics and ultimately how they affect the likelihood of successful coordination in equilibrium.

The common rationale behind Tobin tax-like proposals can be traced back to Pigou’s suggestion to tax actions generating negative externalities; frictions may deter volatile or other harmful investment patterns, the argument goes. However, frictions could backfire if economic performance depends on many aggregates of investment behavior, including the

¹See Tobin (1978).

volume of aggregate investment. A policy maker may worry that the negative effect of frictions on the investment volume may dominate the benefits of reduced volatility. Within the studied model, we dispel such worries. The volume of aggregate investment at the end of the adjustment process (terminal volume henceforth) is invariant to a large class of frictions in a pivotal contingency — we call this the *invariance result*.

Let us illustrate the invariance result on an emerging economy attempting to attract investments and to discourage capital reversals. Exit penalties may help achieve the latter goal, but their effect on the investment volume is seemingly inconclusive. While investors become less likely to exit upon receiving bad news about the economy, they are also less likely to enter in the first place. In our model, these two effects offset each other *exactly*, and under general conditions. The volume of capital that the economy attracts and retains is independent of the frictions. Guided by this invariance result, the policy maker may introduce efficiency-enhancing frictions based on their effect on capital reversals only. Section 2 develops this example further.

The invariance result is driven by relatively weak assumptions. It holds under a large class of information structures that naturally extend those from the static global-games literature. The result is independent of many details of the frictions, and of other details of the investment incentives.

The invariance result provides a reliability test of the static global-game framework. Static global games are a mainstream modelling tool for analyzing coordination processes; see Morris and Shin (2003) for a review of applications. By nature, however, they omit all the dynamic aspects of the studied interactions. Our model distinguishes interactions in which the dynamic elements are important from those where static global games yield reliable predictions. Indeed, there are dynamic specifications of our model where the equilibrium reduces to the equilibrium of a simple static global game. This occurs in dynamic problems where the terminal investment volume is the only determinant of the economic outcome. In such settings, the explicit modelling of the learning process with action adjustments does

not improve the predictions any more than to the static benchmark; frictions do not matter. When economic success, however, depends on several statistics of investors' behavior, as in the example of the emerging economy, the dynamic elements of our model do impact equilibrium behavior. This can be seen in the context of another example.

Consider a dynamic extension of the currency attack model of Morris and Shin (1998). In each period preceding the meeting of a currency board, speculators learn additional information about the fixed state of the economy and adjust their bets on the currency devaluation. Suppose that the board's devaluation decision depends on the aggregate position of the speculators at the time of the meeting, but not on the history or cross-sectional distribution of the positions. As in the example of the emerging economy, the invariance result holds. The terminal aggregate position is independent of the dynamic details of the game. This implies that, although the evolution of the attack may be sensitive to dynamic details, the ex ante equilibrium probability of the attack's success is not. If, however, the history or cross-sectional distribution of the speculators' positions influence the board's decision, dynamic elements become relevant in equilibrium and frictions start to influence the prevalence of the attacks.

Our model is an intended compromise between the tractability of static global games and the richness of dynamic coordination processes. On the one hand, we explicitly model agents' learning and their adjustments to arriving information so that we can analyze the effects of frictions on dynamic investment paths. On the other hand, we do not keep track of flow payoffs, and thus we do not distinguish between early and late economic success. We also model social learning in a reduced form only. In a fully-fledged model of social learning, the agents would receive noisy signals about others' behavior. Since this behavior reflects, in equilibrium, the underlying state of the economy, social learning would lead to an update of agents' beliefs about the state. We simplify the problem by treating arriving information about the state exogenously. Such a reduced-form approach to social learning has been formalized in Dasgupta (2007) and used in Angeletos et al. (2007), Angeletos and Werning

(2006), and Goldstein et al. (2011). Our compromise between the fully dynamic approach and the static equilibrium is in the tradition of the open-loop equilibria in dynamic games (see Fudenberg and Levine 1988). The reduced-form approach to social learning allows us to characterize equilibrium, and thereby to perform a welfare analysis of tax policies in a general environment. Therefore, we view our model as a step towards providing a general testing ground for various Tobin tax-like proposals.

Let us now discuss our contribution in more detail. The invariance result is driven by the assumption that the local properties of the information structure are independent of the realized state; we call this *translational* symmetry. The same assumption underlies the existing static global-game characterizations. For example, selection of the risk-dominant action in Carlsson and van Damme (1993) or selection of the Laplacian² action in Morris and Shin (2003) are driven by this assumption. Kováč and Steiner (2008) use the symmetry to derive a partial equilibrium characterization in a two-stage global game. Although such a symmetry assumption drives uniqueness and characterization results in all global-game models, our invariance result and the analysis in a dynamic setting are novel.³

As a second contribution, we provide sufficient conditions for the existence of monotone equilibria in the dynamic model. The main challenge is to dispense with supermodularity, an assumption which drives existence results in static global games but whose intertemporal analog is overly restrictive.⁴ Our third contribution consists of a characterization result. For settings with fast learning — that is, when the precision of the private information quickly increases across rounds — we show equilibrium existence, provide its characterization, and prove its independence from the assumed error distributions.

The importance of the dynamic aspects of coordination processes has been well recog-

²Morris and Shin use the term “Laplacian action” for the action preferred by an agent who has uniform belief about the aggregate action.

³On its face, translational symmetry is strong, as it requires uninformative prior belief, but it approximately holds for general priors when private information is sufficiently precise; see for instance Frankel et al. (2003). We focus on the uniform prior in the main text, and relegate the extension to general priors to a supplement, available at <http://www.kellogg.northwestern.edu/faculty/steiner/htm/supplement.pdf>.

⁴As discussed by Echenique (2004), the assumption of intertemporal strategic complementarities is very restrictive. See Vives (2009) for a particular approach.

nized. One stream of the literature focuses on dynamic adjustments to an evolving economic environment, see Burdzy et al. (2001), or Chassang (2010). Our paper falls into the class of dynamic models where agents learn about a fixed economic environment. In Chamley (2003), Angeletos and Werning (2006), and Angeletos et al. (2007) agents learn from endogenous public signals such as prices or early coordination outcomes. Since the public signals restore common knowledge in these models, they typically exhibit equilibrium multiplicity. Dasgupta (2007) provides a particular but tractable model of private social learning, within a class of monotone equilibria, equivalent to the exogenous private learning process employed in this paper. Unlike the public learning processes, the private ones preserve strategic uncertainty and equilibrium uniqueness.

The paper is organized as follows. In the next section, we formalize the example of the exit tax in an emerging economy. Section 3 describes the general model and highlights its symmetry property. Section 4 states and proves the invariance result. Section 5 demonstrates existence and characterization results when agents learn fast. Section 6 presents several applied settings in which we use the invariance result to design welfare-enhancing frictions.

2 Example: Sudden Investment Reversals

As an illustration, we design an intuitive welfare improving friction in a simple dynamic coordination game. An emerging economy is opening up to international capital. Foreign investors make decisions in an early and an interim stage of the economic transformation. In the early stage, each investor i of the continuum $[0, 1]$ invests 0 or 2 units in the economy. The early decisions are partially reversible in the interim stage: those who have invested early can withdraw 1 unit, and those who have not invested early can invest 1 unit. We use the same action labels 0 and 1 at each decision node, as in Figure 1.

Investment is risky and costly. Investor i 's terminal payoff before tax is $b^i(o - 1/2)$, where $b^i = a_1^i + a_2^i$ denotes the investor's terminal investment (or bet), $o \in \{0, 1\}$ represents




Figure 1: Dynamic investment problem.

the economic failure or success of the economy, and $1/2$ is the cost of unit investment. The economic outcome o is a function of the aggregate behavior and of the economic fundamental. In this example, the economy succeeds if it attracts enough investment, and if it does not experience too large capital reversal:

$$o = \begin{cases} 1 & \text{if } b - e \geq 1 - \theta, \\ 0 & \text{if } b - e < 1 - \theta, \end{cases} \quad (1)$$

where $b = \int_0^1 b^i di$ is the terminal aggregate investment, $e = \int_0^1 a_1^i (1 - a_2^i) di$ is the volume of capital reversal, and θ is the state of the economy.

We impose tax τ on the detrimental reversals. That is, the after-tax total payoff at path $a^i = (a_1^i, a_2^i)$ is:

$$u(a^i, o) = b^i(o - 1/2) - \tau a_1^i(1 - a_2^i).$$

The tax is levied only when an investor invests in the first round and withdraws in the second.

Intuitively, the tax should (and indeed does in our model) discourage both exit and entry, and thus it should reduce the volume of reversals. However, the effect on the investment volume b is a priori ambiguous. A policy maker may worry that the net effect on the investment volume will be negative and dominate the benefits of reduced capital reversals.

The overall tax effects cannot be understood without detailed modelling of strategic

uncertainty. As in the global-game literature, the uncertainty about opponents' actions is induced by asymmetric information about the state in our model. The state θ is drawn from the uniform distribution on $[-1, 3]$. Investors privately learn about the state as follows. Each investor i receives a private signal $x_t^i = x_{t+1}^i + \eta_t^i$ in each round $t = 1, 2$, with convention $x_3^i = \theta$. The errors η_t^i are uniformly distributed on $[-\sigma, \sigma]$ and are independent across rounds and agents, where $\sigma \in (0, 1/2)$ is a small scaling parameter. Thus, signal quality increases over time. As shown below, the results do not rely on a particular error distribution. Investors do not observe opponents' actions.

We examine threshold equilibria of the following form: (i) there exists an endogenous critical state θ^* such that the project succeeds for $\theta \geq \theta^*$, and fails for $\theta < \theta^*$, and (ii) investors use threshold strategies: an agent at private action history $h \in \{\emptyset, 1, 0\}$ of length $t = 0, 1$ invests in round $t + 1$ if and only if $x_{t+1}^i \geq x_h^*(\tau)$. The private action history h indicates the action taken by the investor up to round t , $\emptyset$ being the null history at the beginning of the game. We have verified that the game has a unique threshold equilibrium for $\tau = 0$, and $1/10$, with critical state $\theta^*(\tau)$.

Each investor knows the critical state θ^* in equilibrium but not the realized θ . Based on her signal, the investor formulates her belief about the event that θ exceeds θ^* , that is, that the economy succeeds. When the realized state is outside of the interval $I(\sigma) = [\theta^* - 2\sigma, \theta^* + 2\sigma]$, the investors know the outcome of the project in each round because they receive signals far enough from θ^* (see Figure 2a). Investors invest and stay for states above the interval, and they never invest for states below. Inside $I(\sigma)$, however, the stochastic learning leads to volatile investment; investment paths 10 and 01 may occur.

Investors' ex ante welfare is (approximately) determined by the critical state. When σ is small, the ex ante probability that the realized state is in $I(\sigma)$ is negligible. Therefore, the ex ante expected utility of an investor is approximately $0 \times \text{Prob}(\theta < \theta^*) + \text{Prob}(\theta > \theta^*) = (3 - \theta^*)/4$. Thus, ex ante welfare is decreasing in the critical state; see Figure 2b for an illustration.




Figure 2: Ex ante Welfare.

Although the critical state is realized with zero probability, the behavior at θ^* determines the equilibrium value of θ^* (as shown in the equation below), and therefore the behavior in this contingency has a large indirect effect on ex-ante welfare. Thus, our analysis focuses on the characterization of the critical behavior.

The condition for success, (1), is met with equality in the critical state:

$$b^* - e^* = 1 - \theta^*,$$

where b^* and e^* are the equilibrium values of investment volume b and reversal volume e in the critical state. They are computed as follows. We choose an arbitrary value of θ^* and find the optimal thresholds x_h^* by solving a simple single agent optimization problem by backward induction. Then, conditioning on the realized state $\theta = \theta^*$, we use the distribution of an investor's signals and the optimal thresholds to compute the probabilities of all investment paths, represented in Figure 3. These probabilities are independent of the value of θ^* chosen at the beginning of the computation. We use the probabilities to compute the expected values of investment and reversals, b^* and e^* .

The effect of the tax on the aggregate behavior in the critical state is depicted in Figure 3. As expected, the tax τ discourages both exit and entry, and thus the critical volume of reversals e^* decreases with the tax. The effect of the tax on the investment volume is less




Figure 3: Probabilities of playing action $a_t \in \{0, 1\}$ in the critical state θ^* , conditional on reaching history $h \in \{\emptyset, 0, 1\}$. (A) $\tau = 0$, and (B) $\tau = 1/10$.

intuitive; it is invariant with respect to the tax; $b^* = 1$ for both tax levels.⁵ The decrease in b^* caused by diminished early entry is exactly offset by the changes in behavior in the interim stage. Thus, the critical state $\theta^* = 1 - b^* + e^* = e^*$ is affected by the exit tax only via the decrease of the reversal volume e^* . As a result of the tax, $\theta^*(0) - \theta^*(1/10) > 0$, and this difference is independent of σ .

The invariance of the critical investment holds under general conditions, which greatly simplifies the design of welfare improving frictions. In this example, the policy maker can focus on frictions' effect on reversals without worrying about investment volume. We conclude that the simple exit tax increases the ex ante probability that the emerging economy succeeds, and thus its expected welfare.

3 The Model

Our model generalizes the example in many directions. The game has arbitrary (finite) number of rounds. We make no distributional assumptions on the learning process. The payoff functions allow for general investment incentives and specifications of transaction costs. Finally, the outcome rule of the investment project may exhibit a broad class of externalities beyond those in the example.

⁵Using Figure 3A, $b^* = (.5 \times .25) + (.5 \times .25) + 2(.5 \times .75) = 1$. Using 3B, $b^* = (.54 \times .27) + (.46 \times .14) + 2(.46 \times .86) = 1$.

3.1 Payoffs

A continuum of agents $i \in [0, 1]$ make binary decisions $a_t^i \in \{0, 1\}$ in rounds $t \in \{1, \dots, T\}$. We write h^i for a private action path $(a_1^i, \dots, a_t^i)$ up to round $t = |h^i|$, denote the initial history at the beginning of the game by $\emptyset$, and often write $z^i \in \{0, 1\}^T$ for terminal paths. When not needed, we omit the index i . We let $h(t')$ be the truncation of h to the first t' elements, and $a(h, t')$ be the t' th action of h . For $h = (a_1, \dots, a_t)$ and $h' = (a'_1, \dots, a'_{t'})$, we let hh' be the path $(a_1, \dots, a_t, a'_1, \dots, a'_{t'})$ of length $t + t'$.

Agent i 's payoff $u(z^i, o)$ depends on her terminal path z^i and on an *outcome* $o \in \{0, 1\}$ interpreted as the failure or success of a common investment project. Since the outcome is binary, we can always express the payoffs as a linear function of o :

$$u(z, o) = b_z \times o - c_z,$$

where $b_z = u(z, 1) - u(z, 0)$ and $c_z = -u(z, 0)$. We interpret b_z as an agent's bet on the success of the project at the terminal path z — the agent receives the amount b_z only if the project succeeds. The parameter c_z is the cost of placing the bet via the path z . In the applications, c_z will include transaction costs alongside the path z .

The outcome of the project is

$$o = \begin{cases} 1 & \text{if } \int_0^1 d_{z^i} di \geq 1 - \theta, \\ 0 & \text{if } \int_0^1 d_{z^i} di < 1 - \theta, \end{cases}$$

where the parameter d_z is the *success contribution* at the terminal path z ; it describes how conducive z is to success. State θ measures the project's propensity to success. Since the parameters b_z , c_z and d_z can depend on the path z in an arbitrary manner, the model accommodates very general dynamic settings. For instance, the example from Section 2 is captured by the success contributions $d_z = b_z - a(z, 1)(1 - a(z, 2))$. Section 6 contains further applications featuring the role of dispersion or volatility of investment. There, we consider

success contributions of the form $d_z = \varphi(b_z) - \lambda v_z$, where $v_z = \sum_{t=2}^T |a(z, t) - a(z, t-1)|$ measures volatility of the path z . The curvature of the function φ determines the impact of the investment dispersion on the project's outcome, and parameter λ determines how harmful volatility is to the outcome.

3.2 Learning Process

The state θ is an unobserved random variable drawn from a uniform distribution on a bounded interval $[\theta_{min}, \theta_{max}]$. Each agent receives a private signal $x_t^i = x_{t+1}^i + \sigma \eta_t^i$ in each round $t = 1, \dots, T$, with convention $x_{T+1}^i = \theta$. Thus, signal x_t^i is a sufficient statistic for the outcome with respect to the private signals up to round t . This specification simplifies the structure of strategies and notation, but is not essential for our main result. The errors η_t^i are independent across agents and rounds and have continuously differentiable density f_t , and distribution F_t with bounded support $[-1/2, 1/2]$. Densities f_t are bounded from below by $\underline{f} > 0$. We abuse terminology by referring both to $\mathbf{x}^i = (x_1^i, \dots, x_T^i)$ and to x_t^i as the type of agent i . The support of θ contains dominance regions: states below $1 - \max_z d_z - T\sigma$ in which all agents in all rounds know that the project fails, and states above $1 - \min_z d_z + T\sigma$ in which all agents know that the project succeeds. Agents do not observe their opponents' actions.

3.3 Strategies and Equilibrium

A pure strategy s is a family of functions $s_h(x_t)$, one for each $h \in \bigcup_{t=0}^{T-1} \{0, 1\}^t$.⁶ An agent following strategy s plays action $a_t = s_h(x_t)$ at each private history h of length $t-1$. Since beliefs at round t are independent of earlier signals, so is the strategy. The independence of the strategy at round t from signals $x_1, \dots, x_{t-1}$ constrains only indifferent types, and these

⁶While the type space of x_t at each round is bounded, we extend the domain of the strategies to the real line $\mathbb{R}$. The purpose is to simplify the upcoming translation arguments. Let $\underline{x}_t$ and $\bar{x}_t$ be the minimal and the maximal signal of the type space at round t . We extend s_h to $\mathbb{R}$ as follows: $s_h(x) = s_h(\bar{x}_{|h|+1})$ for all $x > \bar{x}_{|h|+1}$, and $s_h(x) = s_h(\underline{x}_{|h|+1})$ for all $x < \underline{x}_{|h|+1}$.

have generically zero measure. We say that s is a *threshold strategy* if, for each path h , there exists x_h^* such that $s_h(x) = 1$ for $x \geq x_h^*$, and $s_h(x) = 0$ for $x < x_h^*$.

An *outcome function* $O : \mathbb{R} \rightarrow \{0, 1\}$ specifies the outcome of the project; $o = O(\theta)$. We say that O is a *threshold outcome function* if there exists a *critical state* θ^* such that $O(\theta) = 1$ for $\theta \geq \theta^*$, and $O(\theta) = 0$ for $\theta < \theta^*$.

Strategy s is a *best response* to outcome function O if

$$s_h(x_{|h|+1}) \in \arg \max_a E[V_{ha}(x_{|h|+2}) | x_{|h|+1}],$$

where

$$V_h(x_{|h|+1}) = \max_a E[V_{ha}(x_{|h|+2}) | x_{|h|+1}], \text{ with } x_{T+1} = \theta, \text{ and } V_z(\theta) = b_z O(\theta) - c_z.$$

When O is a threshold outcome function with a critical state θ^* we simply say that s is the best response to θ^* . To avoid ambiguity, we assume that in the case of a tie, agents invest. Then, the best response to any measurable outcome function O is uniquely defined.

Let $z(\mathbf{x}; s)$ be the terminal path that type $\mathbf{x}$ reaches if she follows strategy s . Assume that all agents use the same strategy s and that the Law of Large Numbers applies to the continuous population of agents. Then the aggregate success contribution $\int d_{z(\mathbf{x}^i; s)} di$ in state θ equals the conditional expectation $E[d_{z(\mathbf{x}^i; s)} | \theta]$. The realized state θ determines the outcome of the project both through the conditional distribution of agents' behavior $z(\mathbf{x}^i; s) | \theta$ and through the critical aggregate success contribution $1 - \theta$ required for the success of the project. We say that outcome function O is *generated* by a strategy s if for $\theta \in [\theta_{min}, \theta_{max}]$

$$O(\theta) = \begin{cases} 1 & \text{if } E[d_{z(\mathbf{x}; s)} | \theta] \geq 1 - \theta, \\ 0 & \text{if } E[d_{z(\mathbf{x}; s)} | \theta] < 1 - \theta. \end{cases}$$

For all outcome functions considered below we assume without loss of generality that $O(\theta) =$

1 for $\theta > \theta_{max}$ and $O(\theta) = 0$ for $\theta < \theta_{min}$. We extend the outcome function to all states on the real line for a technical reason; it simplifies the definition of a translation of O .

An *equilibrium* (O, s) is an outcome function O and a strategy s such that s is the best response to O , and O is generated by the strategy s . A pair (O, s) is a *threshold equilibrium* if both O and s are threshold functions.

Our equilibrium concept is a symmetric, pure-strategy Bayes-Nash equilibrium. Since the agents do not observe others' actions, the usual complications with off-equilibrium beliefs, common in dynamic games, do not arise. Our focus on symmetric, pure-strategy equilibria is essentially without loss of generality. All agents optimize against a common outcome function, and thus their best response strategies could differ or use mixed actions only at indifferent types, and these have generically zero measure.

3.4 Translational Symmetry

We conclude the description of the model by highlighting its important symmetry. The joint density of the state and the type, $f(\theta, \mathbf{x})$, is translation invariant:

$$f(\theta, \mathbf{x}) = f(\theta + \delta, \mathbf{x} + \delta \mathbf{e}), \quad (2)$$

where $\mathbf{e} = (1, 1, \dots, 1)$ is the T -dimensional diagonal vector. The translational symmetry is a consequence of the uniform prior and of the additive errors. The symmetry is inherited by the best response function: when $O'(\theta) = O(\theta + \delta)$, the best response s' to O' is the translation of the best response s to O ; $s'_h(x) = s_h(x + \delta)$ for all h . (The translational symmetry of the density f is violated in neighborhoods of the boundary of its support. This however will not play a role in the analysis. This is because the extreme values of θ at the boundaries of the support lie in the dominance regions in which the project either succeeds or fails independently of agents' actions and the agents' decisions are trivial. The analysis is nontrivial only when θ is the intermediate interval, in which case the translational symmetry

of f applies.)

For any variable, we define its *critical value* as its expectation in the critical state θ^* when all agents best-respond to θ^* : for instance, critical success contribution is

$$d^* = \mathbb{E} [d_{z(\mathbf{x};s)} | \theta^*], \text{ where } s \text{ is the best response to } \theta^*.$$

Since both the joint distribution $f(\theta, \mathbf{x})$ and the best response function are translation invariant, the critical value of any variable is independent of θ^* . This implies equilibrium uniqueness within an important class of equilibria:

Lemma 1. *There exists at most one equilibrium (O, s) with a threshold outcome function O . If it exists then the critical state $\theta^* = 1 - d^*$. Moreover, θ^* is independent of σ .*

Proofs omitted in the main text are relegated to Appendix.

4 Invariance of the Critical Investment

This section presents the main insight of the paper — the invariance result. We will show that the terminal volume of aggregate investment in the critical state θ^* depends solely on payoffs received on extreme investment paths. By Lemma 1, the equilibrium is determined by the behavior at θ^* . Thus, though the invariance result applies in only one — the critical — payoff state, it will have strong consequences on the equilibrium and on the ex ante welfare.

Define the *success premium*

$$S = \max_z u(z, 1) - \max_z u(z, 0)$$

as the benefit gained by an informed optimizing agent when the outcome changes from failure to success. Though the agents in the critical state are never perfectly informed about the outcome in our game, the success premium S , defined by optimization under

complete information, happens to characterize the critical aggregate investment. Let s be the best response to θ^* and let

$$b^* = \mathbb{E} [b_{z(\mathbf{x};s)} | \theta^*]$$

denote the critical aggregate investment.

Invariance Result. *The critical aggregate investment is*

$$b^* = S. \tag{3}$$

In particular, the critical investment is invariant to any policy that does not affect the extremal payoffs $\max_z u(z, 1)$, and $\max_z u(z, 0)$.

The term “policy” refers to any change in the payoff parameters b_z or c_z .

The result provides an immediate equilibrium characterization for settings where the outcome depends on the terminal investment volume, but not on other behavioral aggregates, such as volatility or dispersion of investment:

Corollary 1. *Suppose that $d_z = b_z$ for all z , and that an equilibrium with a threshold outcome function exists. Then:*

1. *the critical state $\theta^* = 1 - S$,*
2. *the ex ante probability of success is invariant to any policy that does not affect the extremal payoffs $\max_z u(z, 1)$, and $\max_z u(z, 0)$, and*
3. *the critical state of the dynamic game is identical to the critical state of the static global game in which agents simultaneously choose between the two extremal paths $\arg \max_z u(z, 1)$, and $\arg \max_z u(z, 0)$.*

The corollary provides a robustness check for static global games. Let us illustrate this on the model of currency attacks by Morris and Shin (1998). A currency board makes a devaluation decision at a pre-announced date. Prior to the decision, speculators choose

whether to short sell the currency, based on their private information. The board devaluates the currency if the aggregate short sales exceed a level determined by economic fundamental (specifically $1 - \theta$). The static game of Morris and Shin is a special case of our model with $T = 1$, $b_0 = c_0 = 0$, $b_1 = 1$ and $c_1 \in (0, 1)$.⁷

In our dynamic model with $T > 1$ the agents gradually learn about the state and repeatedly adjust their bets on devaluation. We extend the devaluation rule to the dynamic settings by assuming that the board devaluates the currency if the aggregate short-sales volume *at the time* of the board meeting exceeds $1 - \theta$. Such a rule is captured by our model with $d_z = b_z$. Parameters c_z are the transaction costs incurred on the paths z . See Subsection 6.1 for further formalization of this example.

In this case, the dynamic interaction reduces in equilibrium to the static game. A static model in which agents simultaneously choose only from the two extremal investment paths $\arg \max_z u(z, 1)$, and $\arg \max_z u(z, 0)$ leads to the same equilibrium critical state θ^* as the dynamic model because $\theta^* = 1 - S$ depends only on the extremal payoffs. Thus, modelling the dynamic adjustments explicitly does not improve the performance of the model; the static global-game framework is justified.

Besides modelling considerations, the invariance result has policy implications. The critical state is independent of the transaction costs along all the non-extreme paths. Thus, frictions in the spirit of the Tobin tax that do not impact the payoffs on the extremal, non-volatile paths do not influence equilibrium probability of successful coordination and welfare in this strategic situation.

The invariance result holds in settings with a general prior over θ as agents' posterior beliefs are approximately independent of the prior information when the private signals are sufficiently precise.⁸ Let $s(\phi, \sigma)$ be the best response to θ^* under prior $\phi(\theta)$ when the scaling parameter is σ , and let $b^*(\phi, \sigma) = E_{\phi, \sigma} [b_{z(\mathbf{x}; s(\phi, \sigma))} | \theta^*]$. We prove in the supplement that, as

⁷Morris and Shin allow b_1 and c_1 to depend on the state θ . We discuss such an extension at the end of this section.

⁸See Frankel et al. (2003) for an application of this argument to a large class of static global games.

$\sigma \rightarrow 0$, the critical investment $b^*(\phi, \sigma)$ converges to the success premium S .

Similarly, we conjecture that the invariance result can be extended to settings in which the translational symmetry of the model is violated by dependence of $b_z(\theta)$ and $c_z(\theta)$ on θ . If $b_z(\theta)$ and $c_z(\theta)$ are continuous then they are approximately constant on a small neighborhood of θ^* and the invariance result (approximately) applies for small enough σ .

4.1 Sketch of Proof

We sketch the proof of the invariance result on the example from Section 2. The formal proof is in Appendix.

We fix the value of the critical state θ^* throughout the discussion. Let

$$r = \max_{z'} u(z', o) - u(z, o) \quad (4)$$

be the regret of an investor who has chosen investment path $z = (a_1, a_2)$ when the outcome is o . Abusing notation, let random variable $r(\mathbf{x}, o)$ be an investor's regret when she receives signals $\mathbf{x} = (x_1, x_2)$, follows her optimal investment strategy $s(\mathbf{x})$ against θ^* , and the outcome happens to be o . Conditionally on the realized state being θ , we can compute the distribution of investors' signals. Using this distribution, we define functions $r_o(\theta)$ for $o = 0, 1$ as $\mathbb{E}_{\mathbf{x}} [r(\mathbf{x}, o) \mid \theta]$; it is the expected regret under the optimal strategy s , when the outcome is o , conditional on the true state being θ .

The core of the proof is the observation that the optimal strategy equalizes expected regret in the critical state across success and failure:

$$r_1(\theta^*) = r_0(\theta^*). \quad (5)$$

The invariance result is a corollary of the regret equalization. Rearranging the last equation




Figure 4: A leftward translation of the strategy shifts functions r_0 and r_1 to the left and decreases ex ante expected regret by area A.

gives:

$$\max_{z'} u(z', 1) - \max_{z'} u(z', 0) = \mathbb{E}_{\mathbf{x}} [u(z(\mathbf{x}), 1) - u(z(\mathbf{x}), 0) \mid \theta^*].$$

The left hand side is S (notice that $S = 1$ for both τ in the example from Section 2), and the right-hand side equals $\mathbb{E}_{\mathbf{x}} [b_{z(\mathbf{x})} \mid \theta^*] = b^*$.

The regret equalization result (5) is implied by the translational invariance of the model. Let us prove (5) by contradiction, using Figure 4. When $\theta \ll \theta^*$, investors receive low signals in both rounds, and they know that the project will fail in equilibrium. Therefore, they do not invest. Thus, if the outcome were a success ($o = 1$), as assumed in the definition of r_1 , they would experience substantial regret, because they would wish they had invested in both rounds. Accordingly, $r_1(\theta) = 1$ for sufficiently low θ . As θ increases, an investor is more likely to draw high signals and invest, and thus her expected regret under $o = 1$ decreases until it becomes null. A symmetric argument explains why r_0 is increasing. The ex ante expected regret corresponds to the shaded area in Figure 4. To see this, recall that when $\theta < \theta^*$, the project fails ($o = 0$), and when $\theta > \theta^*$, the project succeeds ($o = 1$). Therefore, for $\theta < \theta^*$, the relevant conditional expected regret is given by r_0 , and for $\theta > \theta^*$, it is given by r_1 .

To establish a contradiction, suppose that the optimal strategy s does not equalize regrets,

as outlined in Figure 4a, where $r_1(\theta^*) > r_0(\theta^*)$. Translating the threshold strategy s to the left, i.e. decreasing all the thresholds x_h^* by the same amount, translates the regret functions r_0 and r_1 to the left because the model exhibits translational symmetry. The translated strategy results in a lower expected regret, since the shaded area in Figure 4b is reduced by the area A . By the definition in (4), expected regret is obtained by subtracting expected utility from a constant. If expected regret can be decreased by translating the strategy, then expected utility can be increased, contradicting the optimality of the original strategy s .

5 Fast Learning

This section presents one of the tractable specifications of the general model. We establish equilibrium existence and characterization under mild restrictions on the payoffs, when the precision of agents' information increases greatly in each round. The online supplement contains an alternative tractable specification that does not restrict the learning process but it imposes stronger payoff restrictions.

To capture fast learning processes, we rescale the errors as

$$x_t^i = x_{t+1}^i + \sigma^t \eta_t^i, \text{ with } x_{T+1}^i = \theta,$$

keeping densities $f_t(\eta_t)$ unaltered, and we examine the limit $\sigma \rightarrow 0$. The essential property of fast learning is that the ratio of signal precisions across two rounds diverges as $\sigma \rightarrow 0$.⁹ We treat the limit of $\sigma \rightarrow 0$ casually in the main text and relegate detailed proofs to Appendix.

We characterize the equilibrium under the assumption that an agent certain of success prefers action 1 while an agent certain of failure prefers the opposite action.

A1: *For all histories h ,*

$$\max_{h'} u(h1h', 1) > \max_{h'} u(h0h', 1),$$

⁹The results of this section hold if $x_t^i = x_{t+1}^i + \sigma^{k(t)} \eta_t^i$ where k is an increasing function.

$$\max_{h'} u(h1h', 0) < \max_{h'} u(h0h', 0),$$

where $h' \in \{0, 1\}^{T-|h|-1}$ is a continuation path.

Consider the best response to a critical state θ^* . As σ vanishes, agents in each round expect to receive information in the next round that is vastly more precise than their current signal. However uncertain about the outcome an agent is in the current round t , she is nearly sure that she will know the outcome and play optimally in the subsequent rounds. Thus, an agent choosing action a at path h expects to receive payoff $b_{ha1\dots 1} - c_{ha1\dots 1}$ if the project succeeds and $-c_{ha0\dots 0}$ if it fails. She invests if and only if she assigns probability at least p_h^* to success, where p_h^* solves the indifference condition

$$(b_{h11\dots 1} - c_{h11\dots 1}) p_h^* - c_{h10\dots 0} (1 - p_h^*) = (b_{h01\dots 1} - c_{h01\dots 1}) p_h^* - c_{h00\dots 0} (1 - p_h^*). \quad (6)$$

Assumption A1 guarantees that p_h^* exists, is unique, and lies in $(0, 1)$.

The equilibrium critical state θ^* is characterized by agents' behavior in the critical state; by Lemma 1:

$$\theta^* = 1 - d^*.$$

To compute the critical success contribution d^* , we analyze the distribution of posterior beliefs at θ^* . Let $q_t(x_t) = \Pr(\theta \geq \theta^* | x_t)$ be the posterior success probability evaluated by type x_t in round t . In the critical state, the posterior beliefs reflect solely the noise in the private signals rather than information about the outcome. As a result, the posteriors are uniformly distributed in the critical state; $q_t(x_t) | \theta^* \sim U[0, 1]$ for any specification of the error distribution.¹⁰

In the critical state, and at each history h , an agent chooses action 1 with probability $\Pr(q_t \geq p_h^* | \theta^*) = 1 - p_h^*$. Additionally, in the limit of fast learning, the posterior beliefs q_t are independent across rounds. Thus, in the limit as $\sigma \rightarrow 0$, the probability that an agent

¹⁰This property has been used in Guimaraes and Morris (2007) and in Steiner (2006). The derivation is as follows. Let $\alpha = x_t - \theta$ be the error and A its c.d.f. Then $\Pr(q_t(x_t) < p | \theta^*) = \Pr(x_t < \theta^* + A^{-1}(p) | \theta^*) = \Pr(\alpha < A^{-1}(p)) = A(A^{-1}(p)) = p$.

reaches a terminal path z is

$$l_z = \prod_{t=1}^T [a(z, t) (1 - p_{z(t-1)}^*) + (1 - a(z, t)) p_{z(t-1)}^*],$$

where, as before, $a(z, t)$ denotes t 'th action on the path z and $z(t-1)$ is the truncation of z to the first $t-1$ rounds. This gives the following limit characterization result:

Proposition 1. *If A1 holds and there exists $\bar{\sigma}$ such that the game has a threshold equilibrium for each $\sigma < \bar{\sigma}$ then $\lim_{\sigma \rightarrow 0} \theta^*(\sigma) = 1 - \sum_z d_z l_z$.*

We supplement the last result with a sufficient condition for the existence of a threshold equilibrium. It requires that investing contributes to success more than not investing, but the restriction is imposed only for the extreme continuation paths:

A2: *For all action histories h , $d_{h1h'} > d_{h0h'}$ where h' is an extreme continuation history $11 \dots 1$ or $00 \dots 0$ of length $T - |h| - 1$.*

The assumption is relatively weak as it leaves the ranking of success contributions on most paths unspecified, thus allowing for considerable modelling freedom.

Proposition 2. *If A1 and A2 hold then there exists $\bar{\sigma} > 0$ such that the game has a unique threshold equilibrium for each $\sigma \in (0, \bar{\sigma}]$.*

Equilibrium uniqueness was established in Lemma 1. Thus it remains to prove existence. To that end, we show in Appendix that the best response s to θ^* generates a non-decreasing expected success contribution $d(\theta) = \mathbb{E}[d_{z(\mathbf{x}, s)} | \theta]$. The main complication is that typical paths in a neighborhood of θ^* are volatile, and thus assumption A2 does not rank their success contributions. The analysis simplifies when learning is fast. Then each agent at any point of the game assigns very high probability to her knowing the outcome in the next round. Hence, when learning is fast, agents believe that their continuation play will be alongside the extreme histories $11 \dots 1$ or $00 \dots 0$. As the proof shows, the incomplete ranking in A2 then suffices to establish monotonicity of $d(\theta)$.

6 Applications

This section presents four specifications of the general model where the project outcome depends on the terminal aggregate investment and possibly on its distribution across investors or on its historical volatility. We characterize the effects of frictions on the ex ante success probability, relying on the invariance result. We focus on the limit of fast learning throughout the section.

Investment is only partially reversible and the investment opportunities decay over time in all four specifications. We formalize this as follows. Each investor decides whether to enter/stay in a project or exit/stay out of the project in each round. An investor who enters in round t^{in} invests $T + 1 - t^{in}$, and withdraws $T + 1 - t^{out}$ if she leaves the project at round t^{out} . Therefore, an agent repeatedly entering and exiting at rounds $t_1^{in} < t_1^{out} < t_2^{in} < t_2^{out} < \dots < t_K^{in} < t_K^{out}$ has committed to the total investment $\sum_{k=1}^K t_k^{out} - t_k^{in}$. We assign label 1 to the action entering/staying in, and 0 to the action exiting/staying out. An agent receives payoff $1/T$ per unit of investment if the project succeeds. Thus investor's total investment (or bet) is

$$b_z = \frac{1}{T} \sum_{t=1}^T a(z, t). \quad (7)$$

This specification allows for an alternative interpretation. An agent decides at the beginning of each round whether to participate in the project in the current round, and her total investment equals the final number of rounds in which she has participated.

To study the effects of frictions, we specify

$$c_z = cb_z + \tau v_z, \text{ where } v_z = \frac{1}{T} \sum_{t=2}^T |a(z, t) - a(z, t-1)|, \quad (8)$$

with $c + \tau < 1$ so that assumption A1 is satisfied throughout this section. In words, the unit cost of investment is c/T with $c \in (0, 1)$, and the agent pays penalty $\tau/T \geq 0$ per entry or exit. Entry at round 1 is not penalized (this is relaxed in Subsection 6.4).

We assume incentives (7) and (8) in all four applications, but we examine different structures of externalities by varying the success contributions d_z . The encompassing specification for all the four applications is

$$d_z = \varphi(b_z) - \lambda v_z.$$

By varying the curvature of the function φ we capture the impact of investment dispersion on the outcome. Parameter λ specifies the impact of volatility.

6.1 Reduction to the Static Game

This application formalizes our discussion of the currency attacks from Section 4. The outcome depends only on the terminal aggregate investment; $d_z = b_z$. Assumption A2 holds, and thus the game has a unique threshold equilibrium for sufficiently small σ . Corollary 1 implies:

Proposition 3. *The critical state $\theta^* = 1 - S = c$.*

In this setting, the outcome of the dynamic game with $T > 1$ reduces to the outcome of the static game with $T = 1$. The dynamic features of the model are unimportant, and the tax is ineffective.

In the subsequent applications, the outcome depends on additional aspects of investment behavior, apart of its terminal volume. Thus, Corollary 1 cannot be applied, and the dynamic model does not reduce to a static game. Yet, the invariance result will be helpful because it implies that the tax affects the equilibrium only via aggregates of behavior other than the terminal investment volume.

6.2 Frictions and Volatility of Investment

We assume that volatility *ceteris paribus* hampers the outcome: the success contributions are $d_z = b_z - \lambda v_z$. When $\lambda < 1/2$, assumption A2 holds and the game has a unique threshold equilibrium for sufficiently small σ .

Using Proposition 1, we can explicitly compute the distribution of the terminal paths z in the critical state, and thus also the critical volatility v^* defined as $E[v_z \mid \theta^*]$. As one would expect, the critical volatility v^* decreases when a small friction τ is introduced. This, combined with the invariance of the critical investment b^* , formalizes Tobin’s intuition:

Proposition 4. *Introduction of a small transaction cost τ increases the ex ante success probability (and thus welfare):*

$$\left. \frac{d}{d\tau} \theta^*(\tau) \right|_{\tau=0} < 0.$$

6.3 Frictions and Concentration of Investment

In this application, we describe the effect of frictions on the cross-sectional distribution of investment across agents. The management and finance literature (e.g. Carlin and Mayer (2003), or Huddart (1993)) has pointed out several channels through which the concentration of investment may affect the outcome of a project. Investment concentrated among few investors may help mitigate the free-riding problem associated with monitoring of the project. On the other hand, dispersed investment is beneficial when delegation from investors to managers fosters success, as small investment levels make such delegation credible.

To capture such effects, we consider success contributions $d_z = \varphi(b_z)$ with φ increasing and convex or concave — in the convex case concentrated investment fosters and in the concave case hampers the success. Assumption A2 holds, so the game has a unique threshold equilibrium when learning is sufficiently fast.

Let us analyze how the frictions affect the distribution of investment across agents in the critical state. As before, the terminal volume of investment is invariant to frictions. The distribution of investment, however, varies with the level of friction. When frictions are low, agents switch actions often and arrive at investment dispersed among many investors. When frictions are high, investment will be concentrated among those investors who, by accident, happened to be optimistic at the beginning of the game and have invested. Because of the high frictions, they tend to stay in the project and arrive at high investment levels. The next

lemma formalizes this using the concept of mean-preserving spread of Rothschild and Stiglitz (1970).

Lemma 2. *Consider $\tau, \tau' \in [0, 1 - c)$, $\tau < \tau'$. The limit distribution of $b_z|\theta^*$ under τ' is a mean-preserving spread of the limit distribution of $b_z|\theta^*$ under τ .*

Thus, the effect of the frictions on the ex ante success probability is unambiguous when φ is convex or concave:

Proposition 5. *1. If concentrated investment ceteris paribus fosters success then frictions increase the ex ante success probability: if φ is convex, then $\theta^*(\tau)$ is decreasing in τ .*

2. If concentrated investment ceteris paribus hampers success then frictions decrease the ex ante success probability: if φ is concave, then $\theta^(\tau)$ is increasing in τ .*

6.4 Long-Horizon Games

Finally, we assume that both the cross-sectional distribution and the volatility of investment influence the project's outcome. Let the success contribution be $d_z = \varphi(b_z) - \lambda v_z$ with φ increasing, twice differentiable, but not necessarily convex or concave. We abandon the simplifying assumption that investment in the first round is not penalized. By convention, we count investment in round 1 as entry; i.e. we set $a(z, 0) = 0$ for all z , and we modify the definition of the volatility to $v_z = \frac{1}{T} \sum_{t=1}^T |a(z, t) - a(z, t-1)|$. When $2\lambda < \varphi'(b)$ for every b , then A2 holds and the game has a unique threshold equilibrium for sufficiently small σ .

We focus on long horizon, $T \rightarrow \infty$. The invariance result applies in this case in a stronger form. First, unlike in the finite-horizon games, the small entry penalty does not distort the critical investment b^* . When investors pay an entry penalty in the first round, the invariance of b_T^* does not hold for finite T because the payoff for the extreme path $11 \dots 1$ varies with τ . However, as T becomes large, the aggregate investment $b_T^* = S_T = 1 - c - \frac{\tau}{T}$ approximates $1 - c$; the first-round entry penalty becomes negligible compared to the overall incentives.

Second, the effect of frictions on the distribution of investment, emphasized in the previous subsection, diminishes for large horizons. When agents adjust to their randomly evolving posteriors in many rounds, the Law of Large Numbers applies and the dispersion in the number of rounds that agents spend in the project becomes negligible. Therefore, the first part of the critical success contribution, $E[\varphi(b_z)|\theta^*]$ converges to $\varphi(E[b_z|\theta^*]) = \varphi(1 - c)$, as $T \rightarrow \infty$, which is independent of the friction τ .

Finally, we analyze investment volatility. We find that the effect of frictions on volatility is large and non-vanishing in the long-horizon games. We explicitly compute the distribution of the terminal paths z in the critical state, and show in Appendix that v_T^* converges to $\frac{2(1-c)c}{1+2\tau}$, as $T \rightarrow \infty$. The next proposition summarizes:

Proposition 6. *Frictions foster the ex ante success probability in long-horizon games:*

$$\lim_{T \rightarrow \infty} \theta^*(T) = 1 - \varphi(1 - c) + \lambda \frac{2(1 - c)c}{1 + 2\tau},$$

where the right hand side decreases in τ .

In the long-horizon games, frictions do not significantly affect volume of aggregate investment, nor its dispersion across investors. Yet, the frictions significantly reduce volatility of investment in the critical state, thus increasing equilibrium ex ante welfare.

7 Conclusion

This paper presents a tractable dynamic global game in which agents privately learn from an exogenous stream of information and repeatedly adjust their actions. The framework is sufficiently rich to allow for the design of welfare-enhancing frictions. The design problem is simplified by the fact that aggregate investment (in a critical contingency) is invariant to a large family of frictions. Thus, a policy maker, using frictions to influence the volatility or concentration of investment, need not worry about compromising the volume of investment.

Relying on this insight, we have characterized the impact of a simple friction on the coordination outcome in various economic situations. Small switching costs foster successful coordination when the economy benefits from a reduction of volatility or from concentrated investment. Frictions are irrelevant when the history and cross-sectional distribution of investment do not impact the economic outcome. When the economy benefits from dispersed investment, switching costs hamper the coordination outcome.

Invariance of critical aggregate investment is driven by the same assumption as the well-known characterization results in static global games — by the translational symmetry of the model. The paper extends our understanding of the consequences of the translational symmetry assumption from static to dynamic settings, thus significantly expanding the range of possible global-games applications.

We have focused on a dynamic extension of a so-called regime change game in which the outcome of the project is binary. Such a setting exhibits a well-defined critical state separating regions of success and failure, which simplified our analysis. Many applications, however, exhibit several degrees of success. Our preliminary results suggest that the invariance result extends to settings with a finite number of outcomes, and thus we believe that the assumption of the binary outcome space is not driving the main result.

One obvious direction for future research is to endogenize the learning process by social learning — that is, by assuming that agents learn from private noisy observations of the early actions of others. The details of social learning would likely alter the equilibrium coordination outcome. Our model, with the exogenous learning process, provides a benchmark for such an exercise.

A Proofs

A.1 Proof of the Result from Section 3

Proof of Lemma 1. Consider an equilibrium with critical state θ^* . Then $E[d_{z(\mathbf{x};s)}|\theta] \geq 1 - \theta$ for all $\theta > \theta^*$, and $E[d_{z(\mathbf{x};s)}|\theta] < 1 - \theta$ for $\theta < \theta^*$. Continuity of $E[d_{z(\mathbf{x};s)}|\theta]$ with respect to θ implies $1 - \theta^* = E[d_{z(\mathbf{x};s)}|\theta^*] = d^*$. The right-hand side is independent of θ^* , as discussed at the end of Section 3.4. Moreover, d^* is also independent of σ as the model is scale invariant: under two values σ and σ' , the best responses to θ^* satisfy $s_h(\theta^* + \sigma\varepsilon; \sigma) = s_h(\theta^* + \sigma'\varepsilon; \sigma')$, and therefore $\Pr_\sigma(z(\mathbf{x}; s(\sigma)) = z' | \theta^*) = \Pr_{\sigma'}(z(\mathbf{x}; s(\sigma')) = z' | \theta^*)$ for any terminal path z' . $\square$

A.2 Proof of the Result from Section 4

Proof of the Invariance Result. We will prove that an agent's expected regret in the critical state θ^* is independent of the outcome: defining regret as

$$R(z, o) = \max_{z'} u(z', o) - u(z, o), \quad (9)$$

we prove that if agents play best-response s to θ^* then the expected regret is equalized across the success and the failure in the critical state θ^* :

$$E[R(z(\mathbf{x}; s), 1) | \theta^*] = E[R(z(\mathbf{x}; s), 0) | \theta^*]. \quad (10)$$

Rearranging (10) immediately gives the invariance result:

$$\max_{z'} u(z', 1) - \max_{z'} u(z', 0) = E[u(z(\mathbf{x}; s), 1) - u(z(\mathbf{x}; s), 0) | \theta^*] = E[b_{z(\mathbf{x}; s)} | \theta^*].$$

The left-hand side is S , which establishes (3), as needed.

The regret equalization (10) is implied by the following: once the agent optimizes her

strategy s against θ^* , her ex ante expected regret cannot be decreased by a change in θ^* :

$$\mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \tilde{\theta}^*})] \geq \mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \theta^*})], \quad (11)$$

where $\mathbf{1}_{\theta \geq \tilde{\theta}^*}$ is the threshold outcome function with the critical state $\tilde{\theta}^*$. Before we derive (11), we first show how it implies (10).

Consider two thresholds θ^* and $\theta^* + \delta$, $\delta > 0$. Since the outcome $O(\theta)$ differs only when $\theta \in [\theta^*, \theta^* + \delta]$,

$$\begin{aligned} \mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \theta^*})] - \mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \theta^* + \delta})] = \\ \int_{\theta^*}^{\theta^* + \delta} \mathbb{E} [R(z(\mathbf{x}; s), 1) - R(z(\mathbf{x}; s), 0) | \theta] \frac{d\theta}{\theta_{max} - \theta_{min}}. \end{aligned}$$

The left-hand side is non-positive for any δ by optimality of θ^* , (11). Since $\mathbb{E} [R(z(\mathbf{x}; s), o) | \theta]$ is continuous in θ , we have proved that

$$0 \geq \mathbb{E} [R(z(\mathbf{x}; s), 1) - R(z(\mathbf{x}; s), 0) | \theta^*].$$

Considering $\delta < 0$ leads to the opposite inequality.

The last step of the proof, the derivation of (11), relies on the translational symmetry of the model, on (2). For any strategy $\hat{s}$ and its translation $\tilde{s}_h(x) = \hat{s}_h(x + \delta)$ the following holds: the distribution of action paths $z(\mathbf{x}; \hat{s}) | \theta$ under symmetric profile $\hat{s}$, conditional on the realization of the fundamental θ equals distribution $z(\mathbf{x}; \tilde{s}) | (\theta - \delta)$ under $\tilde{s}$, conditional on the fundamental being $\theta - \delta$. Thus, for any outcome function $\hat{O}$, and its translation $\tilde{O}(\theta) = \hat{O}(\theta + \delta)$:

$$\mathbb{E} [R(z(\mathbf{x}; \hat{s}), \hat{O}(\theta))] = \mathbb{E} [R(z(\mathbf{x}; \tilde{s}), \tilde{O}(\theta))].$$

Letting s' be the leftward translation of the optimal strategy s , $s'_h(x) = s_h(x + \delta)$, we

have

$$\mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \theta^* + \delta})] = \mathbb{E} [R(z(\mathbf{x}; s'), \mathbf{1}_{\theta \geq \theta^*})] \geq \mathbb{E} [R(z(\mathbf{x}; s), \mathbf{1}_{\theta \geq \theta^*})] .$$

The last inequality holds because s is the best response to θ^* , and, under the definition (9), payoff maximization is trivially equivalent to regret minimization. Comparing the left and right hand sides gives (11). $\square$

A.3 Proofs of the Results from Section 5

Let us review the fast learning specification. Recall $x_t^i = x_{t+1}^i + \sigma^t \eta_t^i$, with $x_{T+1}^i = \theta$. Each error η_t has a continuous density f_t with bounded support, and f_t is bounded from above and below by some positive $\bar{f}$ and $\underline{f}$. We refer to $\varepsilon_t^i = \frac{x_t^i - \theta}{\sigma^t} = \sum_{t'=t}^T \sigma^{t'-t} \eta_{t'}^i$ as *cumulative error* and denote its density by α_t^σ ; it is bounded from above uniformly across all $\sigma \in (0, 1]$ and converges to f_t as $\sigma \rightarrow 0$. Similarly, for $t' < t$, $x_{t'} - x_t = \sum_{\tau=t'}^{t-1} \sigma^\tau \eta_\tau$ and thus there exists $\bar{\alpha}$ such that the conditional density of $x_{t'}|x_t$ is bounded from above by $\frac{\bar{\alpha}}{\sigma^{t'}}$, for all $\sigma \in (0, 1]$.

The following lemma summarizes the heuristic derivation of the optimal strategy from Section 5.

Lemma 3. *If A1 holds then*

1. *there exists $\bar{\sigma}$ such that for all $\sigma < \bar{\sigma}$, the best response to θ^* is a threshold strategy,*
2. *$\lim_{\sigma \rightarrow 0} F_{|h|+1} \left(\frac{x_h^*(\sigma) - \theta^*}{\sigma^{|h|+1}} \right) = p_h^*$, where $p_h^* \in (0, 1)$ solves the indifference condition (6).*

The second statement implies that the threshold type $x_h^*(\sigma)$ assigns probability p_h^* to the success, as $\sigma \rightarrow 0$.

Proof of Lemma 3. Claim 1: Let $\pi_h^\sigma(x_t) = \mathbb{E} [V_{h1}^\sigma(x_{t+1}) - V_{h0}^\sigma(x_{t+1}) | x_t]$ denote the incentive to invest of the type x_t at history h . We will establish single-crossing of $\pi_h^\sigma(x_t)$ for each h and sufficiently small σ .

Type x_t forms beliefs at round t about her signal x_{t+1} in the next round. If $x_{t+1} > \theta^* + \sum_{t'=t+1}^T \sigma^{t'}/2$ (respectively $x_{t+1} < \theta^* - \sum_{t'=t+1}^T \sigma^{t'}/2$), then the agent will be certain at

$t + 1$ that the project will succeed (fail). The probability that the agent will be certain that the project succeeds at $t + 1$, given x_t , is

$$\Pr_{\sigma} \left(x_{t+1} > \theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} \middle| x_t \right) = F_t \left(\frac{x_t - \theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}}{\sigma^t} \right). \quad (12)$$

Let $\underline{x}_t = \theta_{min} - \sum_{t'=t}^T \frac{\sigma^{t'}}{2}$ and $\bar{x}_t = \theta_{max} + \sum_{t'=t}^T \frac{\sigma^{t'}}{2}$ denote the endpoints of the support of x_t . For each t we distinguish three intervals of x_t :¹¹

$$\left[\underline{x}_t, \theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} - \frac{\sigma^t}{2} \right], \left[\theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} - \frac{\sigma^t}{2}, \theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} + \frac{\sigma^t}{2} \right], \left[\theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} + \frac{\sigma^t}{2}, \bar{x}_t \right]. \quad (13)$$

Consider x_t from the third interval. The expression in (12) converges to 1, as $\sigma \rightarrow 0$, uniformly across all x_t from the third interval. Therefore $\pi_h^{\sigma}(x_t)$ converges to $b_{h11\dots1} - c_{h11\dots1} - b_{h01\dots1} + c_{h01\dots1}$, which is positive by A1. Thus $\pi_h^{\sigma}(x_t) > 0$ on the third interval for small enough σ . By a symmetric argument $\pi_h^{\sigma}(x_t) < 0$ on the first interval for small σ .

Next, consider x_t from the middle interval:

$$\begin{aligned} \pi_h^{\sigma}(x_t) &= \Pr_{\sigma} \left(x_{t+1} < \theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} \middle| x_t \right) (-c_{h10\dots0} + c_{h00\dots0}) \\ &\quad + \int_{\theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}}^{\theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}} (V_{h1}(x_{t+1}) - V_{h0}(x_{t+1})) f_t \left(\frac{x_t - x_{t+1}}{\sigma^t} \right) \frac{dx_{t+1}}{\sigma^t} \\ &\quad + \Pr_{\sigma} \left(x_{t+1} > \theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2} \middle| x_t \right) (b_{h11\dots1} - c_{h11\dots1} - b_{h01\dots1} + c_{h01\dots1}), \end{aligned}$$

and letting

$$M(x_{t+1}) = \left(V_{h1}(x_{t+1}) - V_{h0}(x_{t+1}) - \begin{cases} -c_{h10\dots0} + c_{h00\dots0} & \text{if } x_{t+1} < \theta^*, \\ b_{h11\dots1} - c_{h11\dots1} - b_{h01\dots1} + c_{h01\dots1} & \text{if } x_{t+1} > \theta^*, \end{cases} \right)$$

¹¹The endpoints of the intervals are naturally ordered for $\sigma < 1/2$.

we rewrite $\pi_h^\sigma(x_t)$ as

$$\begin{aligned} \pi_h^\sigma(x_t) = & \Pr_\sigma(x_{t+1} < \theta^* | x_t) (-c_{h10\dots 0} + c_{h00\dots 0}) + \int_{\theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}}^{\theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}} M(x_{t+1}) f_t\left(\frac{x_t - x_{t+1}}{\sigma^t}\right) \frac{dx_{t+1}}{\sigma^t} \\ & + \Pr_\sigma(x_{t+1} > \theta^* | x_t) (b_{h11\dots 1} - c_{h11\dots 1} - b_{h01\dots 1} + c_{h01\dots 1}). \end{aligned} \quad (14)$$

The derivative with respect to x_t of the sum of the first and the third terms is

$$[(c_{h10\dots 0} - c_{h00\dots 0}) + (b_{h11\dots 1} - c_{h11\dots 1} - b_{h01\dots 1} + c_{h01\dots 1})] f_t\left(\frac{x_t - \theta^*}{\sigma^t}\right) \frac{1}{\sigma^t},$$

which is positive and of the order $\frac{1}{\sigma^t}$ because the term in the square brackets is positive by A1 and $f_t\left(\frac{x_t - \theta^*}{\sigma^t}\right)$ is bounded from below by a constant $\underline{f}$.

The derivative of the second term is

$$\int_{\theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}}^{\theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}} M(x_{t+1}) f_t'\left(\frac{x_t - x_{t+1}}{\sigma^t}\right) \frac{dx_{t+1}}{\sigma^{2t}}.$$

Term $M(x_{t+1})$ is bounded as payoffs are bounded. The derivative f_t' is bounded as well as error densities are assumed to be continuously differentiable. The whole integral is of the order $\sigma^{t+1} \frac{1}{\sigma^{2t}} = \frac{1}{\sigma^{t-1}}$. Thus, for sufficiently small σ , the sum of derivatives of the first, third and the second term is positive; that is, $\frac{d}{dx_t} \pi_h(x_t)$ is positive on the middle interval in (13).

For small σ , $\pi_h^\sigma(x_t)$ is negative on the first interval, positive on the third interval, increasing on the middle interval and continuous. Thus, the indifference condition $\pi_h^\sigma(x_t) = 0$ has a unique solution $x_h^*(\sigma)$.

Claim 2: Equation (14) implies that, as $\sigma \rightarrow 0$, $\pi_h^\sigma(\theta^* + \sigma^t \varepsilon)$ converges, uniformly across ε to

$$(1 - F_t(\varepsilon)) (-c_{h10\dots 0} + c_{h00\dots 0}) + F_t(\varepsilon) (b_{h11\dots 1} - c_{h11\dots 1} - b_{h01\dots 1} + c_{h01\dots 1}).$$

Hence $F_t\left(\frac{x_h^*(\sigma) - \theta^*}{\sigma^t}\right)$ converges to the solution of the limit indifference condition (6). $\square$

Proof of Proposition 1. By Lemma 1, $\theta^*(\sigma) = 1 - d^*(\sigma)$ where

$$d^*(\sigma) = \sum_{z'} d_{z'} \Pr_{\sigma}(z(\mathbf{x}; s(\sigma)) = z' | \theta^*),$$

and $s(\sigma)$ is the best response to θ^* . Let R_h be the event that the agent reaches path h .

$$\begin{aligned} \Pr_{\sigma}(z(\mathbf{x}; s(\sigma)) = z' | \theta^*) &= \prod_{t=1}^T \left[a(z', t) \Pr_{\sigma}(x_t \geq x_{z'(t-1)}^*(\sigma) | \theta^* \text{ and } R_{z'(t-1)}(\sigma)) \right. \\ &\quad \left. + (1 - a(z', t)) \Pr_{\sigma}(x_t < x_{z'(t-1)}^*(\sigma) | \theta^* \text{ and } R_{z'(t-1)}(\sigma)) \right]. \end{aligned}$$

To show that the last expression converges to l_t , it suffices to prove that for each t and each path h of length $t - 1$:

$$\lim_{\sigma \rightarrow 0} \Pr_{\sigma}(x_t \geq x_h^*(\sigma) | \theta^* \text{ and } R_h(\sigma)) = 1 - p_h^*.$$

We use

$$\Pr_{\sigma}(x_t \geq x_h^*(\sigma) | \theta^* \text{ and } R_h(\sigma)) = \int \Pr_{\sigma}(x_t \geq x_h^*(\sigma) | \theta^* \text{ and } x_{t-1}) g_h(x_{t-1}) dx_{t-1},$$

where $g_h(x_{t-1})$ is the density of x_{t-1} , conditional on $R_h(\sigma)$ and θ^* . We rewrite this further as

$$\int \Pr_{\sigma}(\varepsilon_t \geq \varepsilon_h^*(\sigma) | \varepsilon_{t-1}) \tilde{g}_h(\varepsilon_{t-1}) d\varepsilon_{t-1},$$

where $\tilde{g}_h(\varepsilon_{t-1}) = g_h(\theta^* + \sigma^{t-1} \varepsilon_{t-1}) \sigma^{t-1}$ is the density of ε_{t-1} conditional on $R_h(\sigma)$ and θ^* , and $\varepsilon_h^*(\sigma) = \frac{x_h^*(\sigma) - \theta^*}{\sigma^t}$. The second statement in Lemma 3 implies that all histories are reached with positive, non-vanishing probability in the critical state, as $\sigma \rightarrow 0$. Thus $\tilde{g}_h(\varepsilon_{t-1})$ is bounded. We will show that $\Pr_{\sigma}(\varepsilon_t \geq \varepsilon_h^*(\sigma) | \varepsilon_{t-1})$ converges to $1 - p_h^*$, for each ε_{t-1} . Then by the Dominated Convergence Theorem, the last integral converges to $1 - p_h^*$.

$$\Pr_{\sigma}(\varepsilon_t \geq \varepsilon_h^*(\sigma) | \varepsilon_{t-1}) = \frac{\int_{\varepsilon_h^*(\sigma)}^{\infty} f_{t-1}(\varepsilon_{t-1} - \sigma \varepsilon_t) \alpha_t^{\sigma}(\varepsilon_t) d\varepsilon_t}{\int_{-\infty}^{\infty} f_{t-1}(\varepsilon_{t-1} - \sigma \varepsilon_t) \alpha_t^{\sigma}(\varepsilon_t) d\varepsilon_t},$$

where α_t^{σ} is the unconditional density of ε_t . Additionally, the second statement of Lemma 3 implies that $\lim_{\sigma \rightarrow 0} \varepsilon_h^*(\sigma) = \eta_h^*$, where $F_t(\eta_h^*) = p_h^*$. Since f_{t-1} and α_t^{σ} are bounded, f_{t-1} is continuous, and $\alpha^{\sigma}(\cdot)$ converges to $f_t(\cdot)$:

$$\lim_{\sigma \rightarrow 0} \Pr_{\sigma}(\varepsilon_t \geq \varepsilon_h^*(\sigma) | \varepsilon_{t-1}) = \lim_{\sigma \rightarrow 0} \int_{\varepsilon_h^*(\sigma)}^{\infty} \alpha_t(\varepsilon_t) d\varepsilon_t = 1 - F_t(\eta_h^*) = 1 - p_h^*.$$

□

Proof of Proposition 2. By Lemma 3, the best response to any θ^* is a threshold strategy. Thus, it suffices to prove that there exists $\bar{\sigma}$ such that for any θ^* and all $\sigma \leq \bar{\sigma}$ the best response to θ^* generates a non-decreasing expected success contribution $d_{\sigma}(\theta)$.

Recall the convention $x_{T+1} = \theta$, and let $x_z^* = \theta^*$ for all z . Define auxiliary functions $d_{\sigma}^t(x_t) = \mathbb{E}[d_{z(\mathbf{x}; \tilde{s}^t)} | x_t]$, where the strategy $\tilde{s}^t$ coincides with the best response s to θ^* up to (including) round $t-1$, and specifies action 0 in round t and thereafter; $\tilde{s}_h^t(x) = s_h(x)$ when $|h| < t-1$ and $\tilde{s}_h^t(x) = 0$ when $|h| \geq t-1$. Notice that $d_{\sigma}(x_{T+1}) = d_{\sigma}^{T+1}(x_{T+1})$.

We prove by induction over t that there exists $\bar{\sigma}$ such that for all $\sigma \leq \bar{\sigma}$, and all $t = 1, \dots, T+1$:

$$\frac{d}{dx_t} d_{\sigma}^t(x_t) \geq 0 \text{ for all } x_t \leq \min_{h: |h|=t-1} x_h^*. \quad (15)$$

For $t = T+1$ this is identical to the claim that $\frac{d}{d\theta} d(\theta) \geq 0$ for all $\theta \leq \theta^*$. The proof of the monotonicity above θ^* is symmetric, and we omit it.

Claim (15) holds for $t = 1$ because $d_{\sigma}^1(x_1) = d_{0\dots 0}$. We show that if the claim holds for $t-1$ then it holds for t . Consider first $x_t \leq \min_{h: |h|=t-2} x_h^* - \frac{\sigma^{t-1}}{2}$. Conditional on any x_t from this range, only signals $x_{t-1} \leq \min_{h: |h|=t-2} x_h^*$ have positive probability density in round $t-1$. For such x_{t-1} , $s_h(x_{t-1}) = 0$ for all h of length $t-2$. Thus, for the considered range of x_t , $d_{\sigma}^t(x_t) = \mathbb{E}[d_{\sigma}^{t-1}(x_{t-1}) | x_t]$. Translation invariance of the joint distribution of

signals, (2) implies that, for any function g , $E[g(x_{t-1})|x_t + \delta] = E[g(x_{t-1} + \delta)|x_t]$, and so, $\frac{d}{dx_t} E[g(x_{t-1})|x_t] = E\left[\frac{d}{dx_{t-1}} g(x_{t-1}) \middle| x_t\right]$. Thus,

$$\frac{d}{dx_t} d_\sigma^t(x_t) = E\left[\frac{d}{dx_{t-1}} d_\sigma^{t-1}(x_{t-1}) \middle| x_t\right] \geq 0$$

by the induction hypothesis.

To close the induction step, it remains to prove (15) for t and for

$$x_t \in \left[\min_{h:|h|=t-2} x_h^* - \frac{\sigma^{t-1}}{2}, \min_{h:|h|=t-1} x_h^* \right].$$

For this range,

$$\begin{aligned} d_\sigma^t(x_t) &= \sum_{h:|h|=t-2} \left(\int_{-\infty}^{x_h^*(\sigma)} d_{h00\dots 0} \Pr(R_h|x_{t-1}) f_{t-1} \left(\frac{x_{t-1} - x_t}{\sigma^{t-1}} \right) \frac{dx_{t-1}}{\sigma^{t-1}} + \right. \\ &\quad \left. \int_{x_h^*(\sigma)}^{+\infty} d_{h10\dots 0} \Pr(R_h|x_{t-1}) f_{t-1} \left(\frac{x_{t-1} - x_t}{\sigma^{t-1}} \right) \frac{dx_{t-1}}{\sigma^{t-1}} \right). \end{aligned}$$

Therefore

$$\begin{aligned} \frac{d}{dx_t} d_\sigma^t(x_t) &= \sum_{h:|h|=t-2} (d_{h10\dots 0} - d_{h00\dots 0}) \Pr(R_h|x_h^*) f_{t-1} \left(\frac{x_h^*(\sigma) - x_t}{\sigma^{t-1}} \right) \frac{1}{\sigma^{t-1}} + \\ &\quad \sum_{h:|h|=t-2} \int d_{hs_h(x_t)0\dots 0} f_{t-1} \left(\frac{x_{t-1} - x_t}{\sigma^{t-1}} \right) \frac{d}{dx_{t-1}} \Pr(R_h|x_{t-1}) \frac{dx_{t-1}}{\sigma^{t-1}}. \end{aligned}$$

We show that the first sum on the right-hand side is positive of order $\frac{1}{\sigma^{t-1}}$, and that the second sum is of order $\frac{1}{\sigma^{t-2}}$. Therefore, $\frac{d}{dx_t} d_\sigma^t(x_t) \geq 0$ for small enough σ .

Let us discuss the first sum: $d_{h10\dots 0} - d_{h00\dots 0} > 0$ by A2. The second claim of Lemma 3 implies that $\Pr(R_h|x_h^*) > 0$ for all histories h when σ is sufficiently small. Finally, for at least one path h , $\frac{x_h^*(\sigma) - x_t}{\sigma^{t-1}}$ is in the support of η_{t-1} for the examined interval of x_t , and thus $f_{t-1} \left(\frac{x_h^*(\sigma) - x_t}{\sigma^{t-1}} \right) > \underline{f}$ for at least one path h of length $t - 2$.

To prove that the second sum is of order $\frac{1}{\sigma^{t-2}}$ we show that $\frac{d}{dx_{t-1}} \Pr(R_h|x_{t-1})$ is of order

$\frac{1}{\sigma^{t-2}}$. We establish a bound on $|\Pr(R_h \mid x_{t-1} + \delta) - \Pr(R_h \mid x_{t-1})|$. Let $z_t(\mathbf{x})$ be the first t elements of $z(\mathbf{x}; s)$. Using this notation, $\Pr(R_h \mid x_{t-1}) = \Pr(z_{t-2}(\mathbf{x}) = h \mid x_{t-1})$, and $\Pr(R_h \mid x_{t-1} + \delta) = \Pr(z_{t-2}(\mathbf{x}) = h \mid x_{t-1} + \delta)$. Recalling that the distribution of $\mathbf{x}$ is translation invariant, the last expression equals $\Pr(z_{t-2}(\mathbf{x} + \delta \mathbf{e}) = h \mid x_{t-1})$. Furthermore

$$\begin{aligned} & \left| \Pr(z_{t-2}(\mathbf{x} + \delta \mathbf{e}) = h \mid x_{t-1}) - \Pr(z_{t-2}(\mathbf{x}) = h \mid x_{t-1}) \right| \\ & \leq \Pr(\{\mathbf{x} : \exists \tau < t-1 \text{ such that } s_{h(\tau-1)}(x_\tau + \delta) \neq s_{h(\tau-1)}(x_\tau)\} \mid x_{t-1}) \\ & \leq \sum_{\tau < t-1} \Pr(\{\mathbf{x} : s_{h(\tau-1)}(x_\tau + \delta) \neq s_{h(\tau-1)}(x_\tau)\} \mid x_{t-1}) \leq \sum_{\tau < t-1} \delta \frac{\bar{\alpha}}{\sigma^\tau}, \end{aligned}$$

where we used in the last step that, for $\tau < t-1$, the conditional density of $x_\tau \mid x_{t-1}$ is bounded by $\frac{\bar{\alpha}}{\sigma^\tau}$. The last expression is of order $\frac{\delta}{\sigma^{t-2}}$. $\square$

A.4 Proofs of the Results from Section 6

Proof of Proposition 4. Critical volatility $v^* = v_\emptyset/T$ where $v_\emptyset$ is defined recursively as:¹²

$$\begin{aligned} v_\emptyset &= p_\emptyset^* v_\emptyset + (1 - p_\emptyset^*) v_1, \\ v_h &= p_h^* (a(h, |h|) + v_{h0}) + (1 - p_h^*) (1 - a(h, |h|) + v_{h1}), \text{ when } |h| = 1 \dots T-2, \\ v_h &= p_h^* a(h, |h|) + (1 - p_h^*) (1 - a(h, |h|)), \text{ when } |h| = T-1, \end{aligned}$$

where probabilities p_h^* are given by the indifference condition (6): $p_\emptyset^* = \frac{c+\tau}{1+2\tau}$. For all histories of length $1, \dots, T-2$, $p_h^* = \frac{c+2\tau}{1+2\tau}$ if h ends with 0, and $p_h^* = \frac{c}{1+2\tau}$ for all h ending with 1. Finally, for all histories of length $T-1$, $p_h^* = c + \tau$ if h ends with 0, and $p_h^* = c - \tau$ for all h ending with 1.

Proposition 1 implies that $\theta^* = 1 - b^* + \lambda v_\emptyset/T$. The critical investment b^* is independent of τ by the invariance result, and thus it suffices to prove that $\left. \frac{dv_\emptyset}{d\tau} \right|_{\tau=0} < 0$. We will prove this by induction over the length of the history h .

¹²Variable v_h is the expected number of switches on the continuation path, conditional on agent reaching action history h , and on the state being critical.

First, we let the reader verify that $\frac{dv_h}{d\tau}\big|_{\tau=0} < 0$ when $|h| = T - 1$. Next, consider h of length $1 \dots T - 2$, and assume $\frac{dv_{h'}}{d\tau}\big|_{\tau=0} < 0$ for histories h' of length $|h| + 1$.

$$\frac{dv_h}{d\tau}\bigg|_{\tau=0} = [2a(h, |h|) - 1] \frac{dp_h^*}{d\tau}\bigg|_{\tau=0} + p_h^* \frac{dv_{h0}}{d\tau}\bigg|_{\tau=0} + (1 - p_h^*) \frac{dv_{h1}}{d\tau}\bigg|_{\tau=0} + \frac{dp_h^*}{d\tau}\bigg|_{\tau=0} (v_{h0} - v_{h1}).$$

Using the expressions for p_h^* , it is straightforward to verify that the first summand is negative. The second and the third summands are negative by the induction hypothesis. The fourth summand is zero as $v_{h0} = v_{h1}$ when $\tau = 0$. The last statement holds because the optimal strategy is history-independent when $\tau = 0$.

Finally,

$$\frac{dv_\emptyset}{d\tau}\bigg|_{\tau=0} = p_\emptyset^* \frac{dv_0}{d\tau}\bigg|_{\tau=0} + (1 - p_\emptyset^*) \frac{dv_1}{d\tau}\bigg|_{\tau=0} + \frac{dp_\emptyset^*}{d\tau}\bigg|_{\tau=0} (v_0 - v_1).$$

The first two summands are negative, as we have established that $\frac{dv_h}{d\tau}\big|_{\tau=0} < 0$ for $h \in \{0, 1\}$, and the last summand is again zero when $\tau = 0$. $\square$

Proof of Lemma 2. Let $(a_1, \dots, a_K; p_1, \dots, p_K)$ be a lottery in which outcomes a_k have a probabilities p_k . Recall that lottery L' is a mean-preserving spread of L if there exist lotteries $Z_1, \dots, Z_K$, each with mean a_k , such that L' can be identified with the compound lottery $(Z_1, \dots, Z_K; p_1, \dots, p_K)$. We write $L' \succ L$ if L' is a mean preserving spread of L .

We say that a set of real numbers $\{b_1, \dots, b_K\}$ is a spread of $\{a_1, a_2\}$ if for each $k \in \{1, \dots, K\}$, $b_k \leq a_1$ or $b_k \geq a_2$, the inequality is strict for some k , and $\min_k b_k \leq a_1 < a_2 \leq \max_k b_k$. We omit the proof of the following lemma:

Lemma 4. 1. If $\{b_1, \dots, b_K\}$ is a spread of $\{a_1, a_2\}$, lotteries L and L' have equal means, and supports $\{a_1, a_2\}$ and $\{b_1, \dots, b_K\}$, respectively, then $L' \succ L$.

2. Suppose $L'_k \succ L_k$ for all $k = 1, \dots, K$. Let $L = (L_1, \dots, L_K; p_1, \dots, p_K)$ and $L' = (L'_1, \dots, L'_K; p_1, \dots, p_K)$ be compound lotteries. Then $L' \succ L$.

Fix τ and omit it from the notation for now. Recall that

$$l_h = \prod_{t'=1}^{|h|} [a(h, t') (1 - p_{h(t')}^*) + (1 - a(h, t')) p_{h(t')}^*]$$

is the probability of reaching path h in the critical state, and let L_T be lottery $((b_z)_z; (l_z)_z)$. Define the success premium at path h as

$$S_h = \max_{z: \exists h' \text{ s.t. } z=hh'} u(z, 1) - \max_{z: \exists h' \text{ s.t. } z=hh'} u(z, 0).$$

Note that $S_z = b_z$ for each terminal path z , and $S_\emptyset = S$. Define a sequence of lotteries $L_t = ((S_h)_{h:|h|=t}, (l_h)_{h:|h|=t})$ for $t = 0, \dots, T$. The definition of L_T coincides with the definition from the previous paragraph.

We can rewrite the indifference condition (6) as

$$S_h = (1 - p_h^*) S_{h1} + p_h^* S_{h0},$$

where $1 - p_h^*$ is the probability that agent chooses action 1 at path h , conditional on the state being θ^* . Thus, for each $t = 1, \dots, T$, the lottery L_t can be identified with the compound lottery

$$((Q_h)_{h:|h|=t-1}; (l_h)_{h:|h|=t-1}),$$

where $Q_h = (S_{h0}, S_{h1}; p_h^*, 1 - p_h^*)$ has mean S_h .

We will prove by induction over t that $L_T(\tau') \succ L_T(\tau)$ for each $\tau, \tau' \in [0, 1 - c)$, $\tau' > \tau$. Notice that $L_1(\tau') \succ L_1(\tau)$ because both lotteries have identical means equal to $S_\emptyset = S$, and support of $L_1(\tilde{\tau})$ is $\{S_0(\tilde{\tau}), S_1(\tilde{\tau})\} = \left\{ \frac{(1-c)(T-1)-\tilde{\tau}}{T}, \frac{(1-c)(T-1)+1+\tilde{\tau}}{T} \right\}$, for $\tilde{\tau} \in \{\tau, \tau'\}$. Thus $\{S_0(\tau'), S_1(\tau')\}$ is a spread of $\{S_0(\tau), S_1(\tau)\}$.

From now on, we write variables associated with tax τ' with an apostrophe and variables associated with τ without an apostrophe. For instance, we use $L'_t = L_t(\tau')$ and $L_t = L_t(\tau)$. Assume for induction $L'_{t-1} \succ L_{t-1}$, so that for each h of length $t-1$ there exists a lottery Z_h

with support $\{S'_g\}_{g:|g|=t-1}$ and probabilities $(z_g^h)_{g:|g|=t-1}$, with mean $E[Z_h] = S_h$, such that L'_{t-1} can be identified with the compound lottery $((Z_h)_{h:|h|=t-1}; (l_h)_{h:|h|=t-1})$.

Define a compound lottery $\hat{Q}_h = ((Q'_g)_{g:|g|=t-1}; (z_g^h)_{g:|g|=t-1})$. It is constructed from Z_h by replacing each outcome S'_g by binary lottery $Q'_g = (S'_{g0}, S'_{g1}; p_g^*, 1 - p_g^*)$ with mean S'_g . By construction, the mean of $\hat{Q}_h$ is S_h .

Lotteries L'_t and L_t can be identified with the compound lotteries

$$\left((\hat{Q}_h)_{h:|h|=t-1}; (l_h)_{h:|h|=t-1} \right), \text{ and } \left((Q_h)_{h:|h|=t-1}; (l_h)_{h:|h|=t-1} \right),$$

respectively.

Using the second statement of Lemma 4, $L'_t \succ L_t$ if $\hat{Q}_h \succ Q_h$ for each path h of length $t-1$. For each h , the means of both $\hat{Q}_h$ and Q_h equal S_h and thus, by the first statement of Lemma 4, it suffices to show that the support of $\hat{Q}_h$ is a spread of the support of Q_h . Support of $\hat{Q}_h$ is $\{S'_g\}_{g:|g|=t}$, whereas support of Q_h is $\{S_{h0}, S_{h1}\}$. Let $b_g = \frac{1}{T} \sum_{t'=1}^{|g|} a(g, t')$ be the investment to which the agent has committed at path g . For each path g of length $t < T$

$$S_g(\tau) = \frac{(T-t)(1-c)}{T} + b_g \begin{cases} +\tau/T & \text{if } h \text{ ends with action 1,} \\ -\tau/T & \text{if } h \text{ ends with action 0,} \end{cases}$$

and for g of length T , $S_g(\tau) = b_g$. Fix any path h of length $t-1$. Then for any $\tau, \tau' \in [0, 1-c)$, $\tau < \tau'$, the following holds: $S_g(\tau') \leq S_{h0}(\tau)$ or $S_g(\tau') \geq S_{h1}(\tau)$ for all g of length t , the inequality is strict for at least some g , and $\min_g S_g(\tau') \leq S_{h0}(\tau) < S_{h1}(\tau) \leq \max_g S_g(\tau')$. $\square$

Proof of Proposition 6. Recall that in the critical state θ^* , and in the limit $\sigma \rightarrow 0$, the agent invests at path h with probability $1 - p_h^*$ where p_h^* is the solution of the indifference condition (6). For all histories of length $t < T-1$, $p_h^* = \frac{c+2\tau}{1+2\tau}$ if h ends with action 0, and $p_h^* = \frac{c}{1+2\tau}$ for all h ending with action 1. Since the probability of playing an action at path h only depends on the last action of h , the sequence a_t constitutes in the critical state a Markov chain with

transition matrix

$$Q(a_{t-1}, a_t) = \begin{pmatrix} \frac{c+2\tau}{1+2\tau} & 1 - \frac{c+2\tau}{1+2\tau} \\ \frac{c}{1+2\tau} & 1 - \frac{c}{1+2\tau} \end{pmatrix},$$

and with $a_0 = 0$.

The investment $b_z = \frac{1}{T} \sum_{t=1}^T a_{z(t)}$ is the average action a_t in the first T rounds of a realization z of the Markov chain. For large T , we can apply the Central Limit Theorem for Markov chains (see Kemeny and Snell (1960)) and approximate the distribution of $b_z|\theta^*$ by the normal distribution $N\left(\tilde{b}, \frac{\omega^2}{T}\right)$. The parameter $\tilde{b}$ is the mean of the chain's steady state distribution π that solves $\pi.Q = \pi$. Thus $(\pi_0, \pi_1) = (c, 1 - c)$, and $\tilde{b} = 1 - c$. We report that the variance parameter is $\omega^2 = (1 - c)c(1 + 4\tau)$, and omit the calculation. For any twice differentiable φ ,

$$\mathbb{E}[\varphi(b_z)|\theta^*] = \varphi(1 - c) + \frac{\varphi''(1 - c)\omega^2}{2T} + \mathbf{o}\left(\frac{1}{T^2}\right) \rightarrow \varphi(1 - c),$$

where $\mathbf{o}\left(\frac{1}{T^2}\right)$ is an expression of order of $\frac{1}{T^2}$.

When T is large, we can approximate critical volatility v_T^* applying the stationary distribution π of the ergodic Markov chain: v_T^* converges to $v^* = \pi_0 Q(0, 1) + \pi_1 Q(1, 0) = \frac{2(1-c)c}{1+2\tau}$. Thus, altogether, $\theta^*(T) \rightarrow 1 - \varphi(1 - c) + \lambda \frac{2(1-c)c}{1+2\tau}$, as $T \rightarrow \infty$. $\square$

References

- Angeletos, G., Hellwig, C., Pavan, A., 2007. Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica* 75 (3), 711–756.
- Angeletos, G., Werning, I., 2006. Crises and prices: Information aggregation, multiplicity, and volatility. *The American Economic Review* 96 (5), 1720–1736.
- Burdzy, K., Frankel, D., Pauzner, A., 2001. Fast equilibrium selection by rational players living in a changing world. *Econometrica* 69 (1), 163–189.

- Carlin, W., Mayer, C., 2003. Finance, investment, and growth. *Journal of Financial Economics* 69 (1), 191–226.
- Carlsson, H., van Damme, E., 1993. Global games and equilibrium selection. *Econometrica* 61 (5), 989–1018.
- Chamley, C., 2003. Dynamic speculative attacks. *The American Economic Review* 93 (3), 603–621.
- Chassang, S., 2010. Fear of miscoordination and the robustness of cooperation in dynamic global games with exit. *Econometrica* 78 (3), 973–1006.
- Dasgupta, A., 2007. Coordination and delay in global games. *Journal of Economic Theory* 134 (1), 195–225.
- Echenique, F., 2004. Extensive-form games and strategic complementarities. *Games and Economic Behavior* 46 (2), 348–364.
- Frankel, D. M., Morris, S., Pauzner, A., 2003. Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory* 108 (1), 1–44.
- Fudenberg, D., Levine, D. K., 1988. Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory* 44 (1), 1–18.
- Goldstein, I., Ozdenoren, E., Yuan, K., 2011. Learning and complementarities in speculative attacks. *The Review of Economic Studies* 78 (1), 263.
- Guimaraes, B., Morris, S., 2007. Risk and wealth in a model of self-fulfilling currency attacks. *Journal of Monetary Economics* 54 (8), 2205–2230.
- Huddart, S., 1993. The effect of a large shareholder on corporate value. *Management Science* 39 (11), 114–149.
- Kemeny, J., Snell, J., 1960. *Finite markov chains*. van Nostrand.

- Kováč, E., Steiner, J., 2008. Reversibility in dynamic coordination problems. CERGE-EI Working Papers.
- Morris, S., Shin, H. S., 1998. Unique equilibrium in a model of self-fulfilling currency attacks. *The American Economic Review* 88 (3), 587–97.
- Morris, S., Shin, H. S., 2003. Global games: theory and applications. In: *Advances in Economics and Econometrics, the Eighth World Congress*. Cambridge University Press.
- Rothschild, M., Stiglitz, J., 1970. Increasing risk: I. a definition. *Journal of Economic Theory* 2 (3), 225–243.
- Steiner, J., 2006. Coordination in a mobile world. CERGE-EI Working Papers.
- Tobin, J., 1978. A proposal for international monetary reform. *Eastern Economic Journal* 4 (3/4), 153–159.
- Vives, X., 2009. Strategic complementarity in multi-stage games. *Economic Theory* 40 (1), 151–171.