Tractable Dynamic Global Games and Applications^{*}

Laurent Mathevet^{\dagger}

JAKUB STEINER[‡]

University of Texas at Austin

Northwestern University and CERGE-EI

February 12, 2013

Abstract

We present a family of tractable dynamic global games and its applications. Agents privately learn about a fixed fundamental, and repeatedly adjust their investments while facing frictions. The game exhibits many externalities: payoffs may depend on the volume of investment, on its volatility, and on its concentration. The solution is driven by an invariance result: aggregate investment is (in a pivotal contingency) invariant to a large family of frictions. We use the invariance result to examine how frictions, including those similar to the Tobin tax, affect equilibrium. We identify conditions under which frictions discourage harmful behavior without compromising investment volume.

^{*}Early drafts of this manuscript circulated under the title "Sand in the Wheels: A Dynamic Global-Game Approach". 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 Edinburgh, NYU, HECER in Helsinki, the University of Texas at Austin, the Kansas Theory Workshop, the Midwest theory conference at Vanderbilt, and the conferences "Information and Coordination: Theory and Applications", EEA-ESEM in OSLO, and SED in Cyprus. Rossella Argenziano and Siyang Xiong provided excellent comments as discussants.

[†]lmath@mail.utexas.edu

[‡]j-steiner@kellogg.northwestern.edu

1 Introduction

Many phenomena modeled as global games, be it bank runs, political revolutions, or creditors' panics, are inherently dynamic processes. Yet, the global games commonly applied to these problems are static, simultaneous-move games. This paper develops a tractable dynamic global game with a flexible parametrization. Under some conditions on payoff externalities, the game reduces in equilibrium to a static global game, justifying the static approach used in the applied literature. Under other conditions, however, the dynamic effects prevail. In these cases, adjustment frictions can significantly affect the ex ante probability of successful coordination, impacting welfare.

A continuum of agents interact in a coordination game where expectations of both high and low economic activity can become self-fulfilling. Agents gradually and privately learn about the fixed state of the economy, form expectations about the coordination outcome of the whole economy, and perpetually 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 the state and of various statistics of investors' behavior, including the terminal volume of investment, the volatility of investment, the exit rate, and 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.

Our main technical contribution is the *invariance result*, which characterizes the volume of aggregate investment at the end of the adjustment process in a critical state of the economy. The state of the economy is drawn from the real line, and agents repeatedly learn about the state from private signals. We examine monotone equilibria with a critical state such that the economy succeeds when the realized state exceeds the critical state and fails otherwise. The invariance result shows that aggregate investment in the critical state depends only on a small subset of the model's parameters. It depends only on payoffs along the two extreme investment paths preferred by an agent who knows that the economy fails or succeeds, respectively. The invariance result allows us to solve for the equilibrium critical state, and to analyze the welfare effects of frictions.

The invariance result is driven by an assumption called *translational* symmetry, which requires local properties of the information structure to be independent of the realized state. The same assumption underlies the existing solutions in static global games. For example, selection of the risk-dominant action in Carlsson and van Damme (1993) or selection of the Laplacian¹ action in Morris and Shin (2003) are driven by this assumption. Kováč and Steiner (2013) 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 its application to a dynamic setting are novel.

Let us illustrate the invariance result on an emerging economy attempting to attract foreign 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 (in the critical state). 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 provides a robustness check for the static global-game framework. Our model distinguishes interactions in which the dynamic elements are important from those in which static global games yield reliable predictions. When terminal investment volume is the only determinant of the economic outcome, modelling the dynamics explicitly does not change the predictions, and frictions do not matter. When economic success, however, depends on several statistics of investors' behavior, as in the example of the emerging

 $^{^{1}}$ Morris and Shin use the term "Laplacian action" for the action preferred by an agent who has uniform belief about the aggregate action.

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. In this setting, the invariance result implies that the terminal aggregate position is independent of the dynamic details of the game. 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 positions influence the board's decision, dynamic elements become relevant and frictions start to influence the equilibrium prevalence of the attacks.

Our model is a compromise between the tractability of static global games and the richness of dynamic coordination processes. 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). We briefly study an extension with explicit social learning in Subsection 8.4 and find that the effect of social learning can be expressed as an additive correction to our baseline result. The correction has the natural form of an informational externality and vanishes for some settings.

Let us now discuss our contributions beyond the invariance result. The second contribu-

tion is to 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.² 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; see Proposition 4.

The importance of the dynamic aspects of coordination processes has been well recognized. 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. The next section formalizes the example of the exit tax in an emerging economy. Section 3 describes the general model and highlights its symmetry property. Section 4 states the invariance result. Section 5 demonstrates existence and characterization results when agents learn fast. Section 6 analyzes an alternative tractable setting which does not require the learning process to be fast. Section 7 studies welfareenhancing frictions using the invariance result. Section 8 varies payoffs and information structure in four extensions.

 $^{^{2}}$ As discussed by Echenique (2004), the assumption of intertemporal strategic complementarities is very restrictive. See Vives (2009) for a particular approach.



Figure 1: Dynamic investment problem. The end nodes specify the terminal investments.

2 Example: Investment Reversals

Consider an emerging economy opening up to international capital. Foreign investors make decisions at an early and an interim stage of the economic transformation. In the early stage, each investor $i \in [0, 1]$ invests 0 or 2 units. The early decisions are partially reversible at 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), $o \in \{0, 1\}$ is the failure or success of the economy, and 1/2 is the cost of unit investment. The outcome *o* is a function of the aggregate behavior and of the state of the economy. The economy succeeds if it attracts enough investment, and if it does not experience too large a capital reversal:

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

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

Consider a tax τ on the detrimental reversals, making the payoff for path $a^i = (a_1^i, a_2^i)$:

$$u(a^{i}, o) = b^{i}(o - 1/2) - \tau a_{1}^{i}(1 - a_{2}^{i}).$$

Intuitively, the tax discourages both exit and entry, and thus it reduces the reversal volume. However, the effect on the investment volume b, and hence the total tax effect, are a priori ambiguous.

The paper emphasizes the role of strategic uncertainty—uncertainty about the behavior of other agents. To that end, we study the overall tax effects in a dynamic global game. The state θ is drawn from the uniform distribution on [-1,3]. Investors receive private signals $x_t^i = x_{t+1}^i + \eta_t^i$ 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 scaling parameter. Investors do not observe opponents' actions. We examine threshold equilibria with a critical state θ^* such that the project succeeds for $\theta \ge \theta^*$, and fails for $\theta < \theta^*$.

For simplicity, consider two tax levels, $\tau \in \{0, 1/10\}$. We have verified that the game has a unique threshold equilibrium for both tax values. We define welfare as the ex ante expected equilibrium payoff.

Proposition 1. 1. Aggregate investment b in the critical state is independent of τ .

- 2. Volume of capital reversal e in the critical state decreases in τ .
- 3. Welfare increases in τ , when σ is sufficiently small.

The first claim follows from the invariance result applied to this example. The invariance of the critical investment simplifies the welfare analysis of the tax. The policy maker can focus on the tax effect on reversals without worrying about investment volume.

Let us first derive welfare as a function of the critical state θ^* , and then analyze the equilibrium value of θ^* . The investors form private beliefs about the success of the economy,



Figure 2: (A) Welfare equals $\Pr(\theta \ge \theta^*)$ when the probability of interval $I(\sigma)$ vanishes. (B) Welfare is decreasing in the critical state.

that is, about the event $\theta > \theta^*$. When θ is outside of the interval $I(\sigma) = [\theta^* - 2\sigma, \theta^* + 2\sigma]$, the investors know the outcome in both rounds because they receive signals far enough from θ^* . Investors invest and stay in states above the interval, and they never invest in states below. Inside $I(\sigma)$, stochastic learning leads to heterogeneous volatile investment; investment paths 10 and 01 occur. 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 $0 \times \operatorname{Prob}(\theta < \theta^*) + (2 - 1) \times \operatorname{Prob}(\theta > \theta^*) = (3 - \theta^*)/4$. See Figure 2.

Although the critical state is realized with zero probability, the behavior at θ^* has a large indirect effect on welfare. This is because the equilibrium value of θ^* is determined by the behavior in the critical contingency itself, since the condition for success, (1), must be met with equality in the critical state:

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

where b^* and e^* are the equilibrium investment and reversal volumes b and e in the critical state. They are computed as follows. Take an arbitrary value of θ^* and find the optimal strategy, solving a simple single-agent optimization problem by backward induction. Then, use the conditional distribution of an investor's signals, $(x_1^i, x_2^i) | \theta^*$, and the optimal strategy to compute the probabilities of all investment paths. These probabilities are independent of



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) tax $\tau = 0$, and (B) tax $\tau = 1/10$.

the conjectured value of θ^* due to the translational symmetry of the model; see Subsection 3.4 for details. Finally, use the probabilities to compute the expected values of investment and reversals, b^* and e^* .

The impact of the tax on the critical aggregate behavior is depicted in Figure 3. The tax discourages both exit and entry, and thus e^* decreases with the tax, confirming part 2 of Proposition 1. The effect of the tax on investment is less intuitive: it is invariant with respect to the tax. Indeed, $b^* = 1$ for both tax levels, confirming claim 1.³ The decrease in b^* caused by diminished early entry is exactly offset by the changes in behavior at the interim stage. Thus, the critical state $\theta^*(\tau) = 1 - b^* + e^*(\tau) = e^*(\tau)$ is affected by the exit tax only via the decrease of the reversal volume $e^*(\tau)$. Success probability increases with the tax, as $\theta^*(0) - \theta^*(1/10) > 0$, and this difference is independent of σ .

The invariance of the critical investment holds in a general setting, described in the next section.

3 The Model

We now generalize the example in various ways. The model has an arbitrary finite number of rounds, we make no distributional assumptions on the signal errors, the payoffs allow

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

for general investment incentives and transaction costs, and we allow for a broad class of externalities.

3.1 Payoffs

A continuum of agents $i \in [0, 1]$ make binary decisions $a_t^i \in \{0, 1\}$ in rounds $t \in \{1, \ldots, T\}$. We write h^i for a private action path (a_1^i, \ldots, a_t^i) up to round $t = |h^i|$, denote the initial history at the beginning of the game by \emptyset , and write $z^i \in \{0, 1\}^T$ for terminal paths. When not needed, we omit the index *i*. 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 $h = (a_1, \ldots, a_t)$ and $h' = (a'_1, \ldots, a'_{t'})$, hh' is the path $(a_1, \ldots, a_t, a'_1, \ldots, a'_{t'})$ 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 project. Since the outcome is binary, the payoffs are linear in o without loss of generality:

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

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 project's success—the agent receives the amount b_z only if the project succeeds. The parameter c_z is the cost of placing the bet via path z. In the applications, c_z will include transaction costs alongside path z.

The outcome of the project is

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

where d_z is the success contribution of path z; it describes how conducive z is to success. State θ measures the project's propensity to succeed.

The parameters b_z , c_z and d_z are arbitrary functions of path z, making the model very

general. For example, it accommodates early investments being more productive than later investments, e.g., $d_z = \beta a(z, 1) + a(z, 2)$ with $\beta > 1$. It allows us to analyze the role of investment volatility and investment dispersion: in Section 7, we consider success contributions $d_z = \varphi(b_z) - \lambda v_z$, where $v_z = \sum_{t=2}^{T} |a(z,t) - a(z,t-1)|$ measures the volatility of path z. The curvature of function φ determines the impact of investment dispersion on the outcome.

3.2 Learning Process

The state θ is an unobserved random variable drawn from a uniform distribution on an 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 $t = 1, \ldots, T$, with convention $x_{T+1}^i = \theta$. Thus, x_t^i is a sufficient statistic for the outcome with respect to the private signals up to round t.⁴ 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 $\mathbf{x}^i = (x_1^i, \ldots, x_T^i)$ 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. Thus, the respective extreme paths are dominant in the dominance regions. 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 $h \in \bigcup_{t=0}^{T-1} \{0, 1\}^{t, 5}$ An agent following strategy s plays action $a_t = s_h(x_t)$ at each private history h of length t - 1. The

⁴This specification simplifies the structure of strategies and notation, but is not essential for our main result.

⁵While the 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 \overline{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: $s_h(x) = s_h(\overline{x}_{|h|+1})$ for all $x > \overline{x}_{|h|+1}$, and $s_h(x) = s_h(\underline{x}_{|h|+1})$ for all $x < \underline{x}_{|h|+1}$.

independence of the strategy at round t from signals x_1, \ldots, x_{t-1} is essentially without loss of generality because beliefs (and hence best responses) at round t are independent of earlier signals. We say that s is a threshold strategy if, for each path h, there exists x_h^* such that $s_h(x) = 1$ for $x \ge x_h^*$, and $s_h(x) = 0$ for $x < x_h^*$.

An outcome function $O : \mathbb{R} \longrightarrow \{0, 1\}$ 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 $\theta \ge \theta^*$, and $O(\theta) = 0$ for $\theta < \theta^*$.⁶

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

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

where

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

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

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. Applying the law of large numbers to the continuous population, the aggregate success contribution $\int d_{z(\mathbf{x}^i;s)} di$ in state θ equals the conditional expectation $\mathbb{E}\left[d_{z(\mathbf{x}^i;s)} | \theta\right]$. We say that outcome function O is generated by a strategy s if for $\theta \in [\theta_{min}, \theta_{max}]$

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

⁶For 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 is a pair (O, s) such that s is the best response to O, and O is generated by s. It 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 compromise between the fully dynamic approach and the simultaneous-move model is in the tradition of the open-loop equilibria in dynamic games; see Fudenberg and Levine (1988). The 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 they could use mixed actions, only at indifferent types, which have zero measure.

3.4 Translational Symmetry

We conclude the description of the model by highlighting its translational symmetry. The joint density of the state and the type is translation invariant:

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

where **e** is the *T*-dimensional diagonal vector. The translational symmetry is implied by the uniform prior and 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 f is violated in neighborhoods of the boundary of its support. This, however, will not play a role in the analysis, as the boundaries lie in the dominance regions in which the project either succeeds or fails independently of agents' actions, and the decisions are trivial. The analysis is nontrivial only when θ is in the intermediate interval, where the translational symmetry of fapplies.)

For any variable, define its *critical value* as its expectation in the critical state θ^* when

all agents best-respond to θ^* : for instance, the critical success contribution is

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

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 satisfies $\theta^* = 1 - d^*$. Moreover, θ^* is independent of σ .

Proofs omitted in the main text are in Appendix.

4 Invariance of the Critical Investment

This section presents the main insight of the paper—the invariance result. It states that the terminal volume of aggregate investment in the critical state θ^* depends solely on payoffs received on extreme investment paths. By Lemma 1, equilibrium is determined by the behavior at θ^* . Thus, though the invariance result applies only in the critical state, it has strong indirect consequences on 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, 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

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

denote the critical aggregate investment.

Proposition 2 (Invariance result). The critical aggregate investment satisfies

$$b^* = S. (3)$$

Specifically, 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 a change in the payoff parameters b_z or c_z .

The invariance result is useful only insofar as the threshold equilibrium exists. This section assumes its existence, and Sections 5 and 6 contain two different sets of sufficient conditions for the existence of the threshold equilibrium.

The result directly implies an equilibrium characterization whenever the outcome depends on the terminal investment, but not on other 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, 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, agents gradually learn about the state and repeatedly adjust their bets on devaluation. Suppose that the board devaluates the currency if the aggregate short-sales volume at the time of the board meeting exceeds $1 - \theta$; in our notation, $d_z = b_z$. See Subsection 7.1 for further formalization. A static model in which agents simultaneously choose between short selling without further adjustments and not short selling at all has the same equilibrium critical state θ^* as the dynamic model because $\theta^* = 1 - S$ depends only on the extremal payoffs, and these are identical across the dynamic and the static model. Thus, the static global-game framework is justified when $d_z = b_z$.

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 payoffs on the extremal, non-volatile paths do not influence the equilibrium probability of successful coordination and welfare in this strategic situation.

4.1 Sketch of the Proof

Our proof of the invariance result emphasizes the role of strategic uncertainty. Agents uncertain about the aggregate action find it difficult to coordinate on the efficient action path for fear of regret from miscoordinating with others. The proof of invariance result

⁷Morris and Shin allow b_1 and c_1 to depend on the state θ . We discuss such an extension at the end of Subsection 8.1.

highlights the role of regret and thus naturally extends the arguments supporting the riskdominance criterion. Indeed, the risk dominant action in a symmetric 2×2 game is the one that minimizes the regret from failing to take the best response to the opponent's action.⁸

We sketch the proof of the invariance result on the example from Section 2. The formal proof is in Appendix. Fix the value of the critical state θ^* throughout the subsection. Let

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

be the regret of an investor who has chosen path z when the outcome is o. Abusing notation, let random variable $r(\mathbf{x}, o)$ be an investor's regret when she receives the series of signals \mathbf{x} , follows the best response $s(\mathbf{x})$ against θ^* , and the outcome is o. For $o \in \{0, 1\}$, define $r_o(\theta) = \mathbf{E}_{\mathbf{x}} [r(\mathbf{x}, o) | \theta]$; it is the expected regret under the optimal strategy s, when the outcome is o, conditional on the realized 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:

$$r_1\left(\theta^*\right) = r_0\left(\theta^*\right).\tag{5}$$

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

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

The left-hand side is S, and the right-hand side equals $E_{\mathbf{x}} \left[b_{z(\mathbf{x})} \mid \theta^* \right] = b^*$.

Before proving regret equalization, we discuss Figure 4. When $\theta \ll \theta^*$, investors receive low signals in both rounds, they know that the project will fail, and 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. Accordingly, $r_1(\theta) = 1$ for sufficiently low θ . As θ increases,

⁸We are grateful to the anonymous referee for emphasizing this connection.



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.

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 applies for $r_0(\theta)$. The ex ante expected regret corresponds to the shaded area. 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 $\theta > \theta^*$, by r_1 .

We prove (5) by contradiction, applying the translational symmetry of the model. Suppose that the optimal strategy s does not equalize regrets, as in Figure 4a, where $r_1(\theta^*) > r_0(\theta^*)$. Translation of the threshold strategy s to the left, i.e. reduction of all the thresholds x_h^* by the same amount, shifts the regret functions r_0 and r_1 to the left. 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 has decreased, then expected utility has 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. Section 6 contains an alternative tractable specification that does not restrict the information structure but imposes stronger conditions on payoffs.

We model fast learning by letting $x_t^i = x_{t+1}^i + \sigma^t \eta_t^i$, with $x_{T+1}^i = \theta$, keeping densities $f_t(\eta_t)$ fixed, and we send $\sigma \to 0$. The property of fast learning driving the results of this section is that the ratio of signal precisions across two rounds diverges as $\sigma \to 0$.⁹ We treat the limit of $\sigma \to 0$ casually in the main text and relegate detailed proofs to Appendix.

We prove the existence of a threshold equilibrium under the following two assumptions: A1: For all histories h,

$$\max_{h'} u(h1h', 1) > \max_{h'} u(h0h', 1),$$
$$\max_{h'} u(h1h', 0) < \max_{h'} u(h0h', 0),$$

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

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

Assumption A1 requires that an agent certain of success prefers action 1 while an agent certain of failure prefers 0. Assumption A2 restricts investing to contribute to success more than not investing, but only for the extreme continuation paths. A2 is relatively weak as it leaves the ranking of success contributions on most paths unspecified, thus allowing for considerable modelling freedom.

Proposition 3. There exists $\overline{\sigma} > 0$ such that the game has a unique threshold equilibrium for each $\sigma \in (0, \overline{\sigma}]$.

Lemma 1 has established equilibrium uniqueness. To prove existence, we show in Appendix that the best response s to θ^* generates a non-decreasing expected success contribution $d(\theta) = \mathbb{E}\left[d_{z(\mathbf{x};s)}|\theta\right]$. Thus, there must be a state at which the project transitions from failure to success. The main complication is that typical paths in a neighborhood of θ^* are

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

volatile, and thus assumption A2 does not rank their success contributions. The analysis is simplified when learning is fast. Then each agent at any point of the game assigns very high probability to her knowing the outcome in future rounds. Hence, when learning is fast, agents believe that their continuation play will be alongside the extreme paths 11...1 or 00...0. The incomplete ranking in A2 then suffices to establish the monotonicity of $d(\theta)$.

We now characterize the threshold equilibrium. Consider the best response to a critical state θ^* . As σ vanishes, an agent choosing action a at path h expects to receive payoff $b_{ha1...1} - c_{ha1...1}$ if the project succeeds and $-c_{ha0...0}$ 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

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

Assumption A1 guarantees that $p_h^* \in (0, 1)$.

Under Lemma 1, θ^* is characterized by agents' behavior in the critical state, $\theta^* = 1 - d^*$. To compute the critical success contribution d^* , we analyze the distribution of posterior beliefs at θ^* . Let $q_t(x_t) = \Pr(\theta \ge \theta^* | x_t)$ 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. Guimaraes and Morris (2007) and Steiner (2006) show that 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 $\Pr(q_t \ge p_h^* | \theta^*) = 1 - p_h^*$. Additionally, as $\sigma \to 0$, the posterior beliefs q_t are independent across rounds. Thus, the limit probability that an agent reaches a terminal path z is

$$l_{z} = \prod_{t=1}^{T} \left[a(z,t) \left(1 - p_{z(t-1)}^{*} \right) + (1 - a(z,t)) p_{z(t-1)}^{*} \right],$$
(7)

¹⁰The proof is as follows: let $\alpha = x_t - \theta$ be the error and A its c.d.f. Then $\Pr(q_t(x_t)$

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. In summary:

Proposition 4. Independent of the assumed error distributions, $\lim_{\sigma\to 0} \theta^*(\sigma) = 1 - \sum_z d_z l_z$, with the limit critical probabilities l_z of reaching the terminal path z defined in (6) and (7).

6 General Learning

This section presents a setting with a general learning process, without relying on the limit of fast learning. The setting is a small perturbation of a simple frictionless specification. We will show that the equilibrium of the perturbed game is a small perturbation of the threshold equilibrium of the frictionless game.

The agents face incentives

$$b_z = \sum_{t=1}^{T} a(z, t)$$
, and $c_z = cb_z + \tau \sum_{t=2}^{T} \mathbf{1}_{a(z,t)=r(z(t-1))}$,

with $c \in (0, 1)$ and where a general tax penalizes arbitrarily specified action $r(h) \in \{0, 1\}$ at each path $h \neq \emptyset$. The success contributions are $d_z = b_z - \lambda v_z$, where, unlike before, "volatility" v_z is arbitrarily specified for now. The information structure is the same as in the baseline model, $x_t^i = x_{t+1}^i + \sigma \eta_t^i$, where σ is a fixed constant.

The threshold equilibrium exists when $\tau = \lambda = 0.^{11}$ Under Corollary 1, the critical state $\theta^* = 1 - S = 1 - T(1 - c)$. When τ or λ are positive, the threshold equilibrium may fail to exist, as the best response to θ^* need not be a threshold strategy, or expected success contribution may fail to be monotone when the best response is history dependent. However, a threshold equilibrium exists if the frictions and the impact of volatility are low:

¹¹The best response s to θ^* is a threshold strategy, because an agent invests in each round, if and only if her posterior belief that $\theta \ge \theta^*$ is at least c. The best response s generates a threshold outcome function because expected success contribution $E[d_{z(\mathbf{x};s)}|\theta] = E[b_{z(\mathbf{x};s)}|\theta]$ is non-decreasing in θ .

Proposition 5. There exist $\overline{\tau}$ and $\overline{\lambda}$ such that the game has a unique threshold equilibrium for all $(\tau, \lambda) \in [0, \overline{\tau}] \times [0, \overline{\lambda}]$.

In the remainder of this section we consider a tax system that punishes entries and exits: r(h) = 1 - a(h, |h|) for all histories $h \neq \emptyset$ (recall that a(h, |h|) denotes the last action of history h). We provide a sufficient condition on volatility v_z under which a small tax fosters success. The condition requires that repeating an action from round t - 1 at round t leads to weakly lower volatility than does switching the action, regardless of the continuation path: **A3**: For any path h and continuation path h' such that |h| + 1 + |h'| = T,

$$v_{h\tilde{a}h'} \le v_{ha'h'},\tag{8}$$

where $\tilde{a} = a(h, |h|)$ is the last action of h, and $a' = 1 - \tilde{a}$ is the opposite action.

For instance, v_z can be the number of exits and entries, $\sum_{t=2}^{T} |a_{t-1} - a_t|$, number of "exits", $\sum_{t=2}^{T} \max\{0, a_{t-1} - a_t\}$, number of "entries", $\sum_{t=2}^{T} \max\{0, a_t - a_{t-1}\}$, or v_z can increase non-linearly with the number of entries and exits on the path z.

Assume that $\lambda > 0$ is sufficiently small so that Proposition 5 assures existence of the threshold equilibrium for small enough τ . Then:

Proposition 6. If volatility v_z satisfies A3, then a small tax penalizing exits and entries fosters success:

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

with strict inequality if inequality (8) is strict for some path h.

The proof relies on the invariance result. As tax does not affect the extremal payoffs, the critical investment is independent of τ . Thus it only remains to prove that the critical volatility decreases in τ .

See Subsection 7.2 for a further discussion of frictions and volatility of investment.

7 Applications

This section studies particular settings with partially irreversible investments and decreasing investment opportunities. These can describe an emerging economy that opens a sector to foreign investors, a currency market prone to attack by speculators, or a political revolution. The project outcome depends on the terminal aggregate investment, and possibly on its distribution across investors and its historical volatility. In one application, we study the effect of small frictions similar to a Tobin tax. We focus on the limit of fast learning throughout the section.

Partial irreversibility of investment is the rule rather than the exception, as there is a wedge between the prices at which one can buy and sell capital (Hartman and Hendrickson, 2002). This can be caused by installation costs, transaction costs, or by informational frictions. Arrow (1968) and Abel et al. (1996) propose related arguments that make a case for partial irreversibility. Caballero et al. (1995) present data from U.S. manufacturers that support some degree of irreversibility. Furthermore, investment opportunities are hardly constant over time. In many sectors of emerging economies, such as the construction of infrastructures, investment opportunities decrease over time.

We model partial irreversibility and declining opportunities as follows. In each round, each investor decides whether to enter/stay in the project or exit/stay out of the project. An investor who enters in round t^{in} invests $T + 1 - t^{in}$ units, and withdraws $T + 1 - t^{out}$ units if she exits the project at round t^{out} . Therefore, an agent repeatedly entering and exiting at rounds $t_1^{in} < t_1^{out} < t_2^{in} < t_2^{out} < \cdots < t_K^{in} < t_K^{out}$ has committed to the total investment $\sum_{k=1}^{K} t_k^{out} - t_k^{in}$ at the end of the game. We assign label 1 to the action of entering/staying in, and 0 to the action of exiting/staying out. An agent receives payoff 1/T per unit of investment if the project succeeds. Thus, the investor's total investment is

$$b_z = \frac{1}{T} \sum_{t=1}^{T} a(z, t).$$
(9)

The model can also be interpreted as one where each agent decides in each round whether to participate in the project in the current round, and her total investment is proportional to the number of rounds in which she has participated.

Costs and volatility are defined as

$$c_z = cb_z + \tau v_z$$
, where $v_z = \frac{1}{T} \sum_{t=2}^{T} |a(z,t) - a(z,t-1)|$, (10)

and $c+\tau < 1$, so that A1 holds throughout this section. In words, the unit cost of investment is c/T with $c \in (0, 1)$, and the agent pays penalty $\tau/T \ge 0$ per entry or exit. Entry at round 1 is not penalized (except in Subsection 7.4).

We assume incentives (9) and (10) in all the subsequent applications, and examine different externality structures, with general specification:

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

where φ is increasing. In the context of developing countries, the positive relationship between success and investment volume is supported by the literature on Foreign Direct Investment. Oliva and Rivera-Batiz (2002) and de Mello (1997) provide evidence that FDI is associated with growth and economic success. Parameter λ measures the impact of volatility and the curvature of φ captures the impact of investment dispersion on the outcome. We assume that $\lambda \in [0, 1/2)$ and $2\lambda < \varphi'(b)$ for every b. Then A2 holds and the game has a unique threshold equilibrium in all four applications for sufficiently small σ .

7.1 Reduction to the Static Game

In this application, the outcome depends only on the terminal aggregate investment: $d_z = b_z$. Thus, this scenario formalizes our discussion of the currency attacks from Section 4. Corollary 1 implies: **Proposition 7.** The critical state $\theta^* = 1 - S = c$. Thus, the outcome of the dynamic game is the same as the outcome of the static game in which agents simultaneously choose among the two extreme paths.

In this setting, 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 from 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.

7.2 Frictions and Volatility of Investment

In this application, we examine the stabilizing role of frictions under the assumption that economic turmoils have a negative impact on economic success. In particular, we analyze policies that are reminiscent of Tobin's proposal to "throw some sand in the wheels" of the economy.¹² The rationale behind these policies can be traced back to Pigou's suggestion to tax actions that generate negative externalities. For example, frictions may deter volatile or other harmful investment patterns. However, a policy maker may worry that the negative effect of frictions on the investment volume may dominate the benefits of reduced volatility. Our invariance result dispels such worries: the policy maker may introduce efficiency-enhancing frictions based solely on their effect on capital reversals.

There is ample evidence that investment volatility hampers successful economic outcomes. Lensink and Morrissey (2006) provide empirical evidence that the volatility of FDI has a negative effect on growth in developing countries. The authors propose several explanations for this relationship; one of them is the impact of FDI on R&D. Their study is corroborated by other papers that show that economic volatility in general, to which the volatility of FDI contributes, is detrimental to growth (e.g. Ramey and Ramey, 1995; Mobarak, 2005) and

 $^{^{12}}$ See Tobin (1978).

to welfare (e.g. Pallage and Robe, 2003). The problem is similar for microeconomic units. Froot et al. (1993) and Minton and Schrand (1999) show that cash flow volatility disturbs and reduces capital expenditure, expenses in R&D, and advertising. In view of those studies, we assume that volatility *ceteris paribus* hampers the outcome of the project: $d_z = b_z - \lambda v_z$ with $\lambda > 0$.

Using Proposition 4, we compute the distribution of terminal paths z in the critical state, and thus also the critical volatility v^* defined as $E[v_z | \theta^*]$, as a function of the tax τ . As one would expect, the critical volatility v^* decreases in τ . At the same time, the invariance result implies that the critical investment b^* is unaffected by τ . The next proposition combines these arguments, formalizing Tobin's intuition:

Proposition 8. A small transaction cost τ increases success incidence (and thus welfare):

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

7.3 Frictions and Concentration of Investment

In this application, we focus on the cross-sectional distribution of investment across agents at the end of the game. The management and finance literature (e.g. Carlin and Mayer, 2003; Huddart, 1993) has pointed out several channels through which the concentration of investment may affect the outcome of a project. Investment concentrated among a few investors may help mitigate free-riding 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.

Let us analyze how 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 frictions. When the tax is low, agents switch actions often and arrive at investment dispersed among many investors. When the tax is 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 9.** 1. If concentrated investment ceteris paribus fosters success then frictions increase success incidence: if φ is convex, then $\theta^*(\tau)$ is decreasing in τ .
 - 2. If concentrated investment ceteris paribus hampers success then frictions decrease success incidence: if φ is concave, then $\theta^*(\tau)$ is increasing in τ .

7.4 Long-Horizon Games

Finally, we assume that both the cross-sectional distribution and the volatility of investment influence the 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. We count investment in round 1 as entry; i.e. we set a(z, 0) = 0 for all z, and let $v_z = \frac{1}{T} \sum_{t=1}^{T} |a(z, t) - a(z, t-1)|$.

We study long-horizon games, $T \to \infty$. 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...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.

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 vanishes. 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 $T \to \infty$, which is independent of τ .

Finally, we find that the effect of frictions on volatility is large and non-vanishing in longhorizon games. We explicitly compute the distribution of 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 $T \to \infty$. The next proposition summarizes this:

Proposition 10. Frictions foster success incidence in long-horizon games:

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

where the right-hand side decreases in τ .

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

8 Extensions

This section explores the robustness of the invariance result to assumptions on information and payoffs. Subsection 8.1 proves the invariance result in the presence of public information when the precisions of private signals diverge. Subsection 8.2 has an auxiliary role; it contains an alternative proof of the invariance result for our baseline setting. In Subsection 8.3, we extend the proof to explore the impact of public information with high precision. In Subsection 8.4, we dissociate the effects of social learning from the effects studied in this paper. Subsection 8.5 allows for multiple project outcomes.

8.1 Negligible public information

Frankel et al. (2003) prove that the impact of public information on equilibrium behavior vanishes in static global games as the precision of private signals diverges. This insight extends to our dynamic game.

We model public information as a non-uniform common prior $\phi(\theta)$. Fix the critical state θ^* throughout this subsection and let $s(\phi, \sigma)$ denote the best response to θ^* under the prior ϕ and the scaling of errors σ . Let $b^*(\phi, \sigma) = E_{\phi,\sigma} \left[b_{z(\mathbf{x},s(\phi,\sigma))} \mid \theta^* \right]$ be the critical investment in this environment.

Proposition 11. Suppose the prior density $\phi(\theta)$ has a bounded support and is continuous. The invariance result holds as the private errors vanish: $\lim_{\sigma\to 0} b^*(\phi, \sigma) = S$.

Similarly, we conjecture that the invariance result extends 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 .

8.2 Alternative proof of the invariance result

As an introduction to the next three subsections, we sketch an alternative proof of the invariance result for our baseline setting. Subsequent subsections extend the proof in three directions.

Define

$$U(\tau_{\theta}, \tau_{x}) = \mathbf{E}\left[u\left(z\left(\mathbf{x} + \tau_{x}\mathbf{e}\right), O\left(\theta + \tau_{\theta}\right)\right)\right],$$

where $z(\mathbf{x})$ is the terminal path reached by type \mathbf{x} who follows equilibrium strategy. To understand the definition, first note that U(0,0) is the expected equilibrium payoff. Generally, $U(\tau_{\theta}, \tau_x)$ is the expected payoff of an agent who follows the equilibrium strategy $s(\mathbf{x})$ and the project succeeds when $O(\theta) = 1$ but the joint density of (θ, \mathbf{x}^i) is $f(\theta - \tau_{\theta}, \mathbf{x}^i - \tau_x \mathbf{e})$ instead of $f(\theta, \mathbf{x}^i)$.

We establish the invariance result by examining derivatives of U at point (0,0). First, optimality of the equilibrium strategy implies

$$\frac{\partial}{\partial \tau_x} U\left(\tau_\theta, \tau_x\right) = 0. \tag{11}$$

Second,

$$\frac{\partial}{\partial \tau_{\theta}} U\left(\tau_{\theta}, \tau_{x}\right) = \frac{b^{*}}{\overline{\theta} - \underline{\theta}},\tag{12}$$

because a small rightward translation of the marginal density of θ yields benefit $b_{z(\mathbf{x}^i)}$ to agent *i* when the realized $\theta \approx \theta^*$. Last,

$$\frac{d}{d\tau}U(\tau,\tau) = \frac{S}{\overline{\theta} - \underline{\theta}}.$$
(13)

To see this, consider the translation of the joint density of (θ, \mathbf{x}) by Δ in the direction of the diagonal vector. Such translation increases the probability that an agent receives payoff $\max_z u(z, 1)$ by $\frac{\Delta}{\theta - \underline{\theta}}$, it decreases the probability that an agent receives payoff $\max_z u(z, 0)$ by $\frac{\Delta}{\overline{\theta} - \underline{\theta}}$, and the probability that θ is in the intermediate region I, in which agents miscoordinate, is unmodified; see Figure 2. Overall, the translation leads to a welfare change of $(\max_z u(z, 1) - \max_z u(z, 0)) \frac{\Delta}{\overline{\theta} - \underline{\theta}}$. The combination of the last three equations implies the invariance result.

8.3 Non-negligible public information

Next, we examine public information that is sufficiently precise to impact the equilibrium behavior. We sketch the extension on a well-known static regime-change game for which we establish a new insight. We discuss dynamic games at the end of the subsection.

Consider the static benchmark setting of Angeletos, Hellwig, and Pavan (2007): agents $i \in [0, 1]$ simultaneously choose actions $a^i \in \{0, 1\}$ after they have received private signals $x^i = \theta + \varepsilon^i$, and a public signal $y = \theta + \eta$ with mutually independent and normally distributed errors $\varepsilon^i \sim N(0, \frac{1}{\alpha})$, and $\eta \sim N(0, \frac{1}{\beta})$. Payoff for action 0 is 0, it is 1 - c for action 1 if the joint investment succeeds, and -c if it fails. The investment succeeds if the aggregate action $a = \int_0^1 a^i di \ge 1 - \theta$. Angeletos, Hellwig, and Pavan show that if $\alpha > \frac{\beta^2}{2\pi}$ then the game has a unique equilibrium, which is monotone and symmetric. Below, we consider signal precisions satisfying the last inequality.

The invariance result implies that, in the absence of the public signals, the aggregate action a is S = 1 - c at θ^* . The presence of the public signal y disturbs the translational symmetry, and thus the invariance result does not hold for a fixed value of y. Yet, as argued below, the invariance result holds in a weaker form, after an aggregation across y.

Assume that θ is drawn from a uniform distribution on $[\underline{\theta}, \overline{\theta}]$ so that the joint density of (θ, y, x^i) is $\frac{\phi(\sqrt{\beta}(y-\theta))\phi(\sqrt{\alpha}(x^i-\theta))}{\overline{\theta}-\underline{\theta}}$. Let the support of the prior overlap with the dominance regions and consider high precisions α and β so that the standard errors of the signals are negligible compared to the support of the prior.¹³

Let $\theta^*(y)$ and $x^*(y)$ be the equilibrium critical state and the threshold private signal as a function of the realized public signal y. Let the set $\Theta^* = \{(\theta, y) \in \mathbb{R}^2 : \theta = \theta^*(y)\}$ be the graph of the function $\theta^*(y)$. The invariance result generalizes as follows:

$$\Pr\left[x^{i} \ge x^{*}(y) \mid (\theta, y) \in \Theta^{*}\right] = S = 1 - c.$$

$$(14)$$

¹³We ignore the effects implied by the finite support of the prior distribution. This can be formalized either by sending the signal precisions to infinity or by enlarging the support of the prior to infinity.



Figure 5: Critical state θ^* as a function of the public signal y, for c = 1/3. Fraction $\frac{\alpha}{\beta^2}$ is kept constant, and precisions α and β diverge. The influence of y on θ^* does not disappear in the limit. We show that $\theta^*(y)$ intersects with the diagonal at 1 - S = c.

The left-hand side generalizes the critical investment b^* .¹⁴

The generalized invariance result (14) can be used to find the ex ante success probability when both signals are highly precise: Consider a diverging series of precisions α_k , β_k with a fixed ratio $\frac{\alpha_k}{\beta_k^2} > \frac{1}{2\pi}$ so that the equilibrium is monotone and unique for every k. Figure 5 depicts $\theta_k^*(y)$ as $k \to \infty$ and demonstrates that the impact of the public signal on equilibrium behavior does not vanish in the limit. When the precision of the public signal is high then the marginal probability density of (θ, y) is concentrated in a neighborhood of the diagonal. Hence, for large k,

$$S = \Pr\left[x^{i} \ge x^{*}(y) \mid (\theta, y) \in \Theta_{k}^{*}\right] \approx \Pr\left[x^{i} \ge x^{*}(y) \mid \theta = \theta_{k}^{**}\right],$$

where θ_k^{**} is the solution of $\theta_k^*(y) = y$. Moreover, the condition for the successful outcome must be met with equality in the critical state, and thus $S \approx 1 - \theta_k^{**}$. Therefore, $\lim_k \theta_k^{**} = 1 - S = c$, and the ex ante probability of successful coordination converges to $\frac{\overline{\theta} - c}{\overline{\theta} - \theta}$ as $k \to \infty$.

The proof of the invariance result from the previous subsection extends to this setting as follows. Let $z(x^i, y)$ be the terminal node reached by agent *i*, and let $O(\theta, y)$ be the

¹⁴Note that the conditioning event Θ^* has 0 probability. We define the conditional probability as $\lim_{\delta \to 0} \Pr\left[x^i \ge x^*(y) \mid (\theta, y) \in \Theta^*_{\delta}\right]$, where $\Theta^*_{\delta} = \{(\theta, y) : \exists \delta' \in [0, \delta] \text{ s.t. } (\theta + \delta', y + \delta') \in \Theta^*\}$.

equilibrium outcome function. Define again

$$U(\tau_{\theta}, \tau_{x}) = \mathbb{E}\left[u\left(z\left(x^{i} + \tau_{x}, y + \tau_{\theta}\right), O\left(\theta + \tau_{\theta}, y + \tau_{\theta}\right)\right)\right],$$

where the triple (θ, y, x^i) is the random variable. It is the ex ante expected payoff when the density of (θ, y, x^i) is $f(\theta - \tau_{\theta}, y - \tau_{\theta}, x^i - \tau_x)$.

Equations (11) and (13) from the previous subsection apply in this setting without modification. In the latter case, the introduction of the random public signal together with the uniform marginal distribution of θ reintroduces translational symmetry, and thus (13) follows from the same argument as in the previous subsection. Equation (12) generalizes as

$$\frac{\partial}{\partial \tau_{\theta}} U\left(\tau_{\theta}, \tau_{x}\right) = \frac{1}{\overline{\theta} - \underline{\theta}} \operatorname{E}\left[a^{i}(x^{i}, y) \mid (\theta, y) \in \Theta^{*}\right].$$
(15)

As in the previous subsection, a small translation of the joint density of (θ, y) affects agents when the realized $\theta \approx \theta^*(y)$, but, as $\theta^*(y)$ now depends on the public signal, the conditioning event has been generalized from the singleton $\{\theta^*\}$ to the set Θ^* . The combination of the three equations implies (14).

We have used a static example in this subsection. The restriction to the static-regime change game enabled us to use the existence of monotone equilibrium established by Angeletos, Hellwig, and Pavan. The derivation of (14), however, does not rely on the static setting, and can be extended to dynamic settings if the existence is established by alternative means.

8.4 Social learning

Our baseline model abstracts from social learning. This subsection illustrates in an example that social learning has an equilibrium impact on the coordination outcome, but the effect of social learning can be separated from those studied in this paper. Further, the effect of social learning vanishes in the limit of fast learning.

Consider a specification of our model with T = 2, $b_z = d_z = \frac{a_1 + a_2}{2}$, and $c_z = cb_z$,

 $c \in (0, 1)$. In the benchmark case without social learning, $b^* = S = 1 - c$ and $\theta^* = c$. We now introduce social learning: each agent *i* observes in round 2, apart from her private signal x_2^i , a first-round action a_1^j of a random agent *j* uniformly drawn from the set of agents, with the draws independent across *i*.

A threshold equilibrium consists of the critical state θ^* and five threshold signals: $x^*(\emptyset)$ and $x^*(a_1^i, a_1^j)$ where $(a_1^i, a_1^j) \in \{0, 1\}^2$. Agents in round 1 choose action 1 if $x_1^i \ge x^*(\emptyset)$. They choose action 1 in round 2 if they have played a_1^i , observed a_1^j , and received a signal $x_2^i \ge x^*(a_1^i, a_1^j)$.

Let $z(\mathbf{x}^i, x_1^j)$ be the terminal node reached by agent *i* if both *i* and *j* follow the equilibrium strategy, and let $O(\theta)$ be the equilibrium outcome function. Define

$$U(\tau_{\theta}, \tau_i, \tau_j) = \mathbb{E}\left[u\left(z\left(\mathbf{x}^i + \tau_i \mathbf{e}, x_1^j + \tau_j\right), O\left(\theta + \tau_{\theta}\right)\right)\right].$$

It is the expected payoff of agent *i* when the density of $(\theta, \mathbf{x}^i, x_1^j)$ is $f(\theta - \tau_{\theta}, \mathbf{x}^i - \tau_i \mathbf{e}, x_1^j - \tau_j)$.

Using the same arguments as in Subsection 8.2, and evaluating the derivatives at (0, 0, 0)we get: $\frac{\partial}{\partial \tau_i} U(\tau_{\theta}, \tau_i, \tau_j) = 0$, $\frac{d}{d\tau} U(\tau, \tau, \tau) = \frac{S}{\overline{\theta} - \underline{\theta}}$, and $\frac{\partial}{\partial \tau_{\theta}} U(\tau_{\theta}, \tau_i, \tau_j) = \frac{b^*}{\overline{\theta} - \underline{\theta}}$. Combining all three equations we get:

$$b^* = S - (\overline{\theta} - \underline{\theta}) \frac{\partial}{\partial \tau_j} U(\tau_{\theta}, \tau_i, \tau_j) \,.$$

The invariance result is amended by the second term on the right-hand side that captures the informational externality of social learning.

Detailed exploration of the new term is beyond the scope of this work. In a parallel project of one of the authors, Loeper et al. (2012) explore a setting in which the social learning term dominates the first term. This happens when poorly informed followers observe the actions of highly informed leaders. If, as in Section 5, the late signals are very precise compared to the early signals, the social learning term vanishes. This is because, in the limit of fast learning, the observation of an opponent's early action conveys negligible information compared to one's own late signal and thus the derivative on the right-hand side vanishes.

8.5 Many outcomes

The investment outcome is binary in the baseline setting. In many applications, however, investments may result in various degrees of success.

Consider a static setting with T = 1, $d_z = b_z = z$, $z \in \{0, 1\}$, and $c_z = cb_z$, $c \in (0, 1)$. As before $u(z, o) = b_z o - c_z$, but the outcome o can attain three values:¹⁵

$$o = \begin{cases} 1 & \text{if } b \ge 1 - \theta, \\ \lambda & \text{if } 1 - \theta > b \ge 1/2 - \theta, \\ 0 & \text{if } 1/2 - \theta > b, \end{cases}$$

with $\lambda \in (0, 1)$. This setting can be solved as a standard static global game following Morris and Shin (2003). Our alternative solution method based on an extension of the invariance result has the advantage of being extendible to dynamic settings; see the discussion at the end of the subsection.

In line with Morris and Shin (2003), the game has a unique equilibrium, which is monotone. Thus, there exists $\underline{\theta}^*$ and $\overline{\theta}^*$ such that $O(\theta) = 1$ if $\theta \ge \overline{\theta}^*$, $O(\theta) = \lambda$ if $\overline{\theta}^* > \theta \ge \underline{\theta}^*$, and $O(\theta) = 0$ if $\underline{\theta}^* > \theta$. The invariance result generalizes as:

$$\lambda \underline{b}^* + (1-\lambda)\overline{b}^* = S = 1 - c, \tag{16}$$

where $\underline{b}^* = \mathbf{E}[b_{z(x)} \mid \underline{\theta}^*]$ and $\overline{b}^* = \mathbf{E}[b_{z(x)} \mid \overline{\theta}^*]$.

To derive (16), define $U(\underline{\tau}, \overline{\tau}, \tau_x)$ to be the expected payoff when agents follow a strategy with threshold $x^* - \tau_x$, and two thresholds in the outcome function, $\underline{\theta}^*, \overline{\theta}^*$, are replaced by $\underline{\theta}^* - \underline{\tau}, \overline{\theta}^* - \overline{\tau}$, respectively. As before, U(0, 0, 0) is the equilibrium welfare.

Evaluate all derivatives at point (0, 0, 0). Using the same arguments as in the previous subsections, $\frac{\partial}{\partial_{\tau_x}}U(\underline{\tau}, \overline{\tau}, \tau_x) = 0$, and $\frac{d}{d\tau}U(\tau, \tau, \tau) = \frac{S}{\overline{\theta}-\underline{\theta}}$. Moreover, $\frac{\partial}{\partial_{\underline{\tau}}}U(\underline{\tau}, \overline{\tau}, \tau_x) = \frac{\lambda \underline{b}^*}{\overline{\theta}-\underline{\theta}}$

¹⁵Extension to arbitrarily many finite values is immediate.

because a small leftward shift in $\underline{\theta}^*$ yields benefits λ to all agents who have invested in contingencies $\theta \approx \underline{\theta}^*$. Similarly, $\frac{\partial}{\partial \tau} U(\underline{\tau}, \overline{\tau}, \tau_x) = \frac{(1-\lambda)\overline{b}^*}{\overline{\theta}-\underline{\theta}}$. The combination of all four equations of this paragraph implies (16).

The criticality condition generalizes as $\underline{b}^* = 1/2 - \underline{\theta}^*$ and $\overline{b}^* = 1 - \overline{\theta}^*$. Moreover, $\underline{\theta}^*$ and $\overline{\theta}^*$ converge to the same limit as $\sigma \to 0$, which we denote by θ^* .¹⁶ From now on we examine the limit $\sigma \to 0$ (we abuse notation by not distinguishing $\underline{b}^*(\sigma)$, $\overline{b}^*(\sigma)$ from their limits). The limit criticality conditions are

$$\underline{b}^* = 1/2 - \theta^*,\tag{17}$$

$$\overline{b}^* = 1 - \theta^*. \tag{18}$$

The system of three linear equations (16), (17), and (18) for three unknowns, \underline{b}^* , \overline{b}^* , and θ^* implies $\theta^* = c - \lambda/2$.

Equations (16), (17), and (18) hold unmodified in threshold equilibria of dynamic settings with many rounds. Thus, when the threshold equilibrium exists, the method extends to our dynamic model.

9 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 ¹⁶To see this, notice that $\underline{\theta}^*, \overline{\theta}^* \in (x^* - \sigma, x^* + \sigma)$. 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 wellknown 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 game applications.

A Proofs

A.1 Proof for Section 3

Proof of Lemma 1. Consider an equilibrium with critical state θ^* . Then $\mathbb{E}\left[d_{z(\mathbf{x};s)} \middle| \theta\right] \ge 1-\theta$ for all $\theta > \theta^*$, and $\mathbb{E}\left[d_{z(\mathbf{x};s)} \middle| \theta\right] < 1-\theta$ for $\theta < \theta^*$. Continuity of $\mathbb{E}\left[d_{z(\mathbf{x};s)} \middle| \theta\right]$ with respect to θ implies $1-\theta^* = \mathbb{E}\left[d_{z(\mathbf{x};s)} \middle| \theta^*\right] = d^*$. 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 $\Pr_{\sigma}(z(\mathbf{x}; s(\sigma)) = z' \middle| \theta^*) = \Pr_{\sigma'}(z(\mathbf{x}; s(\sigma')) = z' \middle| \theta^*)$ for any terminal path z'.

A.2 Proof for 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

$$R(z, o) = \max_{z'} u(z', o) - u(z, o),$$
(19)

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^* :

$$\mathbf{E}\left[R(z(\mathbf{x};s),1)|\theta^*\right] = \mathbf{E}\left[R(z(\mathbf{x};s),0)|\theta^*\right].$$
(20)

Rearranging (20) immediately gives the invariance result:

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

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

The regret equalization (20) 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^* :

$$\mathbf{E}\left[R\left(z(\mathbf{x};s),\mathbf{1}_{\theta\geq\tilde{\theta}^*}\right)\right]\geq\mathbf{E}\left[R\left(z(\mathbf{x};s),\mathbf{1}_{\theta\geq\theta^*}\right)\right],\tag{21}$$

where $\mathbf{1}_{\theta \geq \tilde{\theta}^*}$ is the threshold outcome function with the critical state $\tilde{\theta}^*$. Before we derive (21), we first show how it implies (20).

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

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

The left-hand side is non-positive for any δ by optimality of θ^* , (21). Since $\mathbb{E}\left[R(z(\mathbf{x};s),o) \mid \theta\right]$ is continuous in θ , we have proved that

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

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

The last step of the proof, the derivation of (21), 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 $z(\mathbf{x}; \hat{s}) \mid \theta$ under symmetric profile \hat{s} , conditional on the realization of the fundamental θ equals distribution $z(\mathbf{x}; \tilde{s}) \mid (\theta - \delta)$ under \tilde{s} , conditional on the fundamental being $\theta - \delta$. Thus, for any outcome function \hat{O} , and its translation $\tilde{O}(\theta) = \hat{O}(\theta + \delta)$:

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

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

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

The last inequality holds because s is the best response to θ^* , and, under the definition (19), payoff maximization is trivially equivalent to regret minimization. Comparing the left- and right-hand sides gives (21).

A.3 Proofs for 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 \overline{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 $\sigma \to 0$. Similarly, for t' < t, $x_{t'} - x_t = \sum_{\tau=t'}^{t-1} \sigma^{\tau} \eta_{\tau}$ and thus there exists

 $\overline{\alpha}$ such that the conditional density of $x_{t'}|x_t$ is bounded from above by $\frac{\overline{\alpha}}{\sigma^{t'}}$, 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 $\overline{\sigma}$ such that for all $\sigma < \overline{\sigma}$, the best response to θ^* is a threshold strategy,

2.
$$\lim_{\sigma \to 0} F_{|h|+1}\left(\frac{x_h^*(\sigma) - \theta^*}{\sigma^{|h|+1}}\right) = p_h^*$$
, 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 $\sigma \to 0$.

Proof of Lemma 3. Claim 1: Let $\pi_h^{\sigma}(x_t) = \mathbb{E}\left[V_{h1}^{\sigma}(x_{t+1}) - V_{h0}^{\sigma}(x_{t+1})|x_t\right]$ 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

$$\Pr_{\sigma}\left(x_{t+1} > \theta^* + \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} \middle| x_t\right) = F_t\left(\frac{x_t - \theta^* - \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2}}{\sigma^t}\right).$$
(22)

Let $\underline{x}_t = \theta_{min} - \sum_{t'=t}^T \frac{\sigma^{t'}}{2}$ and $\overline{x}_t = \theta_{max} + \sum_{t'=t}^T \frac{\sigma^{t'}}{2}$ denote the endpoints of the support of x_t . For each t we distinguish three intervals of x_t :¹⁷

$$\left[\underline{x}_{t}, \theta^{*} + \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} - \frac{\sigma^{t}}{2}\right], \quad \left[\theta^{*} + \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} - \frac{\sigma^{t}}{2}, \theta^{*} - \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} + \frac{\sigma^{t}}{2}\right], \quad \left[\theta^{*} - \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} + \frac{\sigma^{t}}{2}, \overline{x}_{t}\right]$$
(23)

Consider x_t from the third interval. The expression in (22) converges to 1, as $\sigma \to 0$, uniformly across all x_t from the third interval. Therefore $\pi_h^{\sigma}(x_t)$ converges to $b_{h11...1}$ –

 $^{^{17} \}mathrm{The}$ endpoints of the intervals are naturally ordered for $\sigma < 1/2.$

 $c_{h11...1} - b_{h01...1} + c_{h01...1}$, 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:

$$\pi_{h}^{\sigma}(x_{t}) = \Pr_{\sigma} \left(x_{t+1} < \theta^{*} - \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} \middle| x_{t} \right) \left(-c_{h10...0} + c_{h00...0} \right) \\ + \int_{\theta^{*} - \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2}}^{\theta^{*} + \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2}} \left(V_{h1}(x_{t+1}) - V_{h0}(x_{t+1}) \right) f_{t} \left(\frac{x_{t} - x_{t+1}}{\sigma^{t}} \right) \frac{dx_{t+1}}{\sigma^{t}} \\ + \Pr_{\sigma} \left(x_{t+1} > \theta^{*} + \sum_{t'=t+1}^{T} \frac{\sigma^{t'}}{2} \middle| x_{t} \right) \left(b_{h11...1} - c_{h11...1} - b_{h01...1} + c_{h01...1} \right),$$

and letting

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

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

$$\pi_h^{\sigma}(x_t) = \Pr_{\sigma}\left(x_{t+1} < \theta^* | x_t\right) \left(-c_{h10\dots0} + c_{h00\dots0}\right) + \int_{\theta^* - \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}}^{\theta^* + \sum_{t'=t+1}^T \frac{\sigma^{t'}}{2}} M(x_{t+1}) f_t\left(\frac{x_t - x_{t+1}}{\sigma^t}\right) \frac{dx_{t+1}}{\sigma^t}$$

$$+\Pr_{\sigma}\left(x_{t+1} > \theta^* | x_t\right) \left(b_{h11\dots 1} - c_{h11\dots 1} - b_{h01\dots 1} + c_{h01\dots 1}\right).$$
(24)

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

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

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 <u>f</u>.

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 (23).

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 (24) implies that, as $\sigma \to 0$, $\pi_h^{\sigma}(\theta^* + \sigma^t \varepsilon)$ converges, uniformly across ε to

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

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

Proof of Proposition 3. By Lemma 3, the best response to any θ^* is a threshold strategy. Thus, it suffices to prove that there exists $\overline{\sigma}$ such that for any θ^* and all $\sigma \leq \overline{\sigma}$ 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 $d_{\sigma}^t(x_t) = \mathbb{E}\left[d_{z(\mathbf{x};\tilde{s}^t)} \mid x_t\right]$, 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 $|h| \ge t-1$. Notice that $d_{\sigma}(x_{T+1}) = d_{\sigma}^{T+1}(x_{T+1})$.

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

$$\frac{d}{dx_t} d^t_{\sigma}(x_t) \ge 0 \text{ for all } x_t \le \min_{h:|h|=t-1} x_h^*.$$
(25)

For t = T + 1 this is identical to the claim that $\frac{d}{d\theta}d(\theta) \ge 0$ for all $\theta \le \theta^*$. The proof of the

monotonicity above θ^* is symmetric, and we omit it.

Claim (25) holds for t = 1 because $d_{\sigma}^{1}(x_{1}) = d_{0...0}$. 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} , $d_{\sigma}^{t}(x_{t}) = \mathbb{E}\left[d_{\sigma}^{t-1}(x_{t-1})|x_{t}\right]$. Translation invariance of the joint distribution of signals, (2) implies that, for any function g, $\mathbb{E}\left[g\left(x_{t-1}\right)|x_{t}+\delta\right] = \mathbb{E}\left[g\left(x_{t-1}+\delta\right)|x_{t}\right]$, and so, $\frac{d}{dx_{t}}\mathbb{E}\left[g(x_{t-1})|x_{t}\right] = \mathbb{E}\left[\frac{d}{dx_{t-1}}g(x_{t-1})|x_{t}\right]$. Thus,

$$\frac{d}{dx_t}d^t_{\sigma}(x_t) = \mathbb{E}\left[\left.\frac{d}{dx_{t-1}}d^{t-1}_{\sigma}(x_{t-1})\right|x_t\right] \ge 0$$

by the induction hypothesis.

To close the induction step, it remains to prove (25) 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,

$$d_{\sigma}^{t}(x_{t}) = \sum_{h:|h|=t-2} \left(\int_{-\infty}^{x_{h}^{*}(\sigma)} d_{h00...0} \Pr(R_{h}|x_{t-1}) f_{t-1} \left(\frac{x_{t-1} - x_{t}}{\sigma^{t-1}} \right) \frac{dx_{t-1}}{\sigma^{t-1}} + \int_{x_{h}^{*}(\sigma)}^{+\infty} d_{h10...0} \Pr(R_{h}|x_{t-1}) f_{t-1} \left(\frac{x_{t-1} - x_{t}}{\sigma^{t-1}} \right) \frac{dx_{t-1}}{\sigma^{t-1}} \right).$$

Therefore

$$\frac{d}{dx_{t}}d_{\sigma}^{t}(x_{t}) = \sum_{h:|h|=t-2} (d_{h10\dots0} - d_{h00\dots0}) \Pr(R_{h}|x_{h}^{*}) f_{t-1} \left(\frac{x_{h}^{*}(\sigma) - x_{t}}{\sigma^{t-1}}\right) \frac{1}{\sigma^{t-1}} + \sum_{h:|h|=t-2} \int d_{hs_{h}(x_{t})0\dots0} f_{t-1} \left(\frac{x_{t-1} - x_{t}}{\sigma^{t-1}}\right) \frac{d}{dx_{t-1}} \Pr(R_{h}|x_{t-1}) \frac{dx_{t-1}}{\sigma^{t-1}}$$

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, $\frac{d}{dx_t}d_{\sigma}^t(x_t) \ge 0$ for small enough σ .

Let us discuss the first sum: $d_{h10...0} - d_{h00...0} > 0$ 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 $|\Pr(R_h | x_{t-1} + \delta) - \Pr(R_h | x_{t-1})|$. Let $z_t(\mathbf{x})$ be the first t elements of $z(\mathbf{x}; s)$. Using this notation, $\Pr(R_h | x_{t-1}) = \Pr(z_{t-2}(\mathbf{x}) = h | x_{t-1})$, and $\Pr(R_h | x_{t-1} + \delta) = \Pr(z_{t-2}(\mathbf{x}) = h | x_{t-1} + \delta)$. Recalling that the distribution of \mathbf{x} is translation invariant, the last expression equals $\Pr(z_{t-2}(\mathbf{x} + \delta \mathbf{e}) = h | x_{t-1})$. Furthermore

$$\left| \Pr\left(z_{t-2}(\mathbf{x} + \delta \mathbf{e}) = h \mid x_{t-1} \right) - \Pr\left(z_{t-2}(\mathbf{x}) = h \mid x_{t-1} \right) \right|$$

$$\leq \Pr\left(\left\{ \mathbf{x} : \exists \tau < t - 1 \text{ such that } s_{h(\tau-1)}(x_{\tau} + \delta) \neq s_{h(\tau-1)}(x_{\tau}) \right\} \middle| x_{t-1} \right)$$

$$\leq \sum_{\tau < t-1} \Pr\left(\left\{ \mathbf{x} : s_{h(\tau-1)}(x_{\tau} + \delta) \neq s_{h(\tau-1)}(x_{\tau}) \right\} \middle| x_{t-1} \right) \leq \sum_{\tau < t-1} \delta \frac{\overline{\alpha}}{\sigma^{\tau}},$$

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 $\frac{\overline{\alpha}}{\sigma^{\tau}}$. The last expression is of order $\frac{\delta}{\sigma^{t-2}}$.

Proof of Proposition 4. By Lemma 1, $\theta^*(\sigma) = 1 - d^*(\sigma)$ where

$$d^*(\sigma) = \sum_{z'} d_{z'} \operatorname{Pr}_{\sigma} \left(z(\mathbf{x}; s(\sigma)) = z' | \theta^* \right),$$

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

$$\Pr_{\sigma}\left(z(\mathbf{x}; s(\sigma)) = z' | \theta^*\right) = \prod_{t=1}^{T} \left[a(z', t) \Pr_{\sigma}\left(x_t \ge x^*_{z'(t-1)}(\sigma) \middle| \theta^* \text{ and } R_{z'(t-1)}(\sigma) \right) + (1 - a(z', t)) \Pr_{\sigma}\left(x_t < x^*_{z'(t-1)}(\sigma) \middle| \theta^* \text{ and } R_{z'(t-1)}(\sigma) \right) \right].$$

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:

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

We use

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

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

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

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 $\sigma \to 0$. Thus $\tilde{g}_h(\varepsilon_{t-1})$ is bounded. We will show that $\Pr_{\sigma}(\varepsilon_t \ge \varepsilon_h^*(\sigma) | \varepsilon_{t-1})$ converges to $1 - p_h^*$, for each ε_{t-1} . Then by the Dominated Convergence Theorem, the last integral converges to $1 - p_h^*$.

$$\Pr_{\sigma} \Big(\varepsilon_{t} \geq \varepsilon_{h}^{*}(\sigma) | \varepsilon_{t-1} \Big) = \frac{\int_{\varepsilon_{h}^{*}(\sigma)}^{\infty} f_{t-1} \left(\varepsilon_{t-1} - \sigma \varepsilon_{t} \right) \alpha_{t}^{\sigma}(\varepsilon_{t}) d\varepsilon_{t}}{\int_{-\infty}^{\infty} f_{t-1} \left(\varepsilon_{t-1} - \sigma \varepsilon_{t} \right) \alpha_{t}^{\sigma}(\varepsilon_{t}) d\varepsilon_{t}},$$

where α_t^{σ} is the unconditional density of ε_t . Additionally, the second statement of Lemma 3 implies that $\lim_{\sigma\to 0} \varepsilon_h^*(\sigma) = \eta_h^*$, 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 $f_t(\cdot)$:

$$\lim_{\sigma \to 0} \Pr_{\sigma} \left(\left| \varepsilon_{t} \geq \varepsilon_{h}^{*}(\sigma) \right| \varepsilon_{t-1} \right) = \lim_{\sigma \to 0} \int_{\varepsilon_{h}^{*}(\sigma)}^{\infty} \alpha_{t}(\varepsilon_{t}) d\varepsilon_{t} = 1 - F_{t} \left(\eta_{h}^{*} \right) = 1 - p_{h}^{*}.$$

A.4 Proofs for Section 6

Let us extend the existing notation: $z(h, \mathbf{x}; s)$ denotes the terminal path that type \mathbf{x} reaches if she starts at path h and follows strategy s in the subsequent rounds. It is defined recursively as $z(h, \mathbf{x}; s) = z(hs_h(x_{|h|+1}), \mathbf{x}; s)$, and $z(h', \mathbf{x}; s) = h'$, if |h'| = T.

Proof of Proposition 5. The proof is divided into Lemmas 4 and 5. The first lemma establishes monotonicity of optimal strategies, and the second lemma proves monotonicity of the outcome function.

Lemma 4. There exists $\overline{\tau}$ such that for all $\tau < \overline{\tau}$ the best response to a threshold outcome function is a threshold strategy.

Proof of Lemma 4. Recall that $\pi_h(x_t; \tau) = \mathbb{E}\left[V_{h1}(x_{t+1}; \tau) - V_{h0}(x_{t+1}; \tau) | x_t\right]$ denotes the incentive to invest. Notice that $\pi_h(x; \tau)$ converges to $\pi_h(x; 0)$ as $\tau \to 0$, uniformly across x, because $|V_h(x_t; 0) - V_h(x_t; \tau)| \leq T\tau$. The function $\pi_h(x; 0) = \Pr(\theta \geq \theta^* | x) - c$ has a unique root $x_h^*(0)$, and $\pi_h(x; 0)$ is bounded away from 0 apart from any δ -neighborhood of $x_h^*(0)$. Thus for each $\delta > 0$, equation $\pi_h(x; \tau) = 0$ has a solution and each solution $x_h^*(\tau)$ lies in the δ -neighborhood of $x_h^*(0)$, for sufficiently small τ . It remains to prove that the solution is unique.

We prove below that $\frac{\partial}{\partial x}\pi_h(x;\tau)$ converges to $\frac{\partial}{\partial x}\pi_h(x;0)$ as $\tau \to 0$, uniformly across x. Since $\frac{\partial}{\partial x}\pi_h(x;0)$ is positive in a sufficiently small neighborhood of $x_h^*(0)$, function $\pi_h(x;\tau)$ satisfies single-crossing, because $\frac{\partial}{\partial x}\pi_h(x;\tau)|_{x=x^*(\tau)} > 0$ for small enough τ . Therefore the best response to θ^* is a threshold strategy, as needed.

To prove that $\lim_{\tau\to 0} \frac{\partial}{\partial x} \pi_h(x;\tau) = \frac{\partial}{\partial x} \pi_h(x;0)$ uniformly across x, we show that for h of length t-1, $\lim_{\tau\to 0} \frac{\partial}{\partial x_t} \mathbb{E}[V_{ha}(x_{t+1};\tau)|x_t] = \frac{\partial}{\partial x_t} \mathbb{E}[V_{ha}(x_{t+1};0)|x_t]$, uniformly across x_t , for both a. For this, we use translation invariance of the joint distribution $f(\theta, \mathbf{x})$; (2). It implies that, for any function $g(x_{t+1})$, $\mathbb{E}[g(x_{t+1})|x_t+\delta] = \mathbb{E}[g(x_{t+1}+\delta)|x_t]$, and thus $\frac{\partial}{\partial x_t} \operatorname{E}\left[g\left(x_{t+1}\right) | x_t\right] = \operatorname{E}\left[\frac{\partial}{\partial x_{t+1}} g\left(x_{t+1}\right) | x_t\right].$ Using the last identity several times,

$$\frac{\partial}{\partial_{x_t}} \mathbb{E}\left[V_{ha}(x_{t+1};\tau) | x_t\right] = \mathbb{E}\left[\left.\frac{\partial}{\partial_{x_{t+1}}} V_{ha}(x_{t+1};\tau) \right| x_t\right] = \dots = \mathbb{E}\left[\left.\frac{\partial}{\partial_{x_T}} V_{z(h,\mathbf{x};s(\tau))}(x_T;\tau) \right| x_t\right],$$

and, similarly, $\frac{\partial}{\partial x_t} \mathbb{E}\left[V_{ha}(x_{t+1}; 0) | x_t\right] = \mathbb{E}\left[\frac{\partial}{\partial x_T} V_{z(h, \mathbf{x}; s(0))}(x_T; 0) | x_t\right]$. Finally, $\frac{\partial}{\partial x_T} V_{z(h, \mathbf{x}; s(\tau))}(x_T; \tau) = f_T\left(\frac{x_T - \theta^*}{\sigma}\right) \frac{1}{\sigma} b_{z(h, \mathbf{x}; s(\tau))}$ and the previous paragraph established that the set of \mathbf{x} on which strategies $s(\tau)$ and s(0) differ, vanishes as $\tau \to 0$. Therefore $\mathbb{E}\left[\frac{\partial}{\partial x_T} V_{z(h, \mathbf{x}; s(\tau))}(x_T; \tau) | x_t\right]$ converges to $\mathbb{E}\left[\frac{\partial}{\partial x_T} V_{z(h, \mathbf{x}; s(0))}(x_T; 0) | x_t\right]$, uniformly across x_t .

Lemma 5. There exists $\overline{\tau} > 0$ and $\overline{\lambda} > 0$ such that for all $(\tau, \lambda) \in [0, \overline{\tau}] \times [0, \overline{\lambda}]$ the best response $s(\tau)$ to a threshold outcome function O generates another threshold outcome function O'.

Proof of Lemma 5. Fix a critical state θ^* and consider function $d(\theta; \tau, \lambda) = \mathbb{E} \left[d_{z(\mathbf{x};s(\tau))} \middle| \theta \right]$ generated by the best response $s(\tau)$ to θ^* . We show that $\frac{d}{d\theta} d(\theta; \tau, \lambda) > -1$ for small enough τ and λ , which establishes single-crossing for $d(\theta; \tau, \lambda) - (1 - \theta)$ with respect to θ .

Let us write $d(\theta; \tau, \lambda)$ as $b(\theta; \tau) - \lambda v(\theta; \tau)$ where $b(\theta; \tau) = \sum_{z'} b_{z'} \Pr(z(\mathbf{x}; s(\tau)) = z' | \theta)$ and $v(\theta; \tau) = \sum_{z'} v_{z'} \Pr(z(\mathbf{x}; s(\tau)) = z' | \theta)$. Thus $\frac{d}{d\theta} d(\theta; \tau, \lambda) = \frac{d}{d\theta} b(\theta; \tau) - \lambda \frac{d}{d\theta} v(\theta; \tau)$. The first summand converges to $\frac{d}{d\theta} b(\theta; 0)$ as $\tau \to 0$ because, as shown in Lemma 4, thresholds $x_h^*(\tau) \to x_h^*(0)$. In the limit case, $\frac{d}{d\theta} b(\theta; 0) = \sum_{t=1}^T \frac{d}{d\theta} \Pr(x_t \ge x_t^* | \theta)$, which is non-negative. Moreover, $\frac{d}{d\theta} v(\theta; \tau)$ is bounded from above, so the right-hand side exceeds -1 for sufficiently small τ and λ .

To finish the proof of Proposition 5, we verify that the outcome function O with threshold $\theta^* = 1 - d^*$, and the best response s to O constitute a threshold equilibrium. Indeed, by Lemma 4, s is a threshold strategy. By Lemma 5, s generates a threshold outcome function O'. But O' = O as needed because we chose $\theta^* = 1 - d^*$.

Proof of Proposition 6. It suffices to analyze derivative $\frac{d}{d\tau}v^*(\tau)$. All the derivatives in the proof are evaluated at $\tau = 0$.

We divide the proof into Lemmas 6 and 7. The first lemma states that the tax discourages agents from exits and entries.

Lemma 6. For all histories h ending with action 0 or 1, $\frac{d}{d\tau}x_h^*(\tau) > 0$ or $\frac{d}{d\tau}x_h^*(\tau) < 0$, respectively.

Proof of Lemma 6. We apply the Implicit Function Theorem on the indifference condition $\pi_h(x^*(\tau);\tau) = 0$:

$$\frac{\partial \pi_h(x;0)}{\partial x}\bigg|_{x=x_h^*(0)}\frac{dx_h^*(\tau)}{d\tau} + \frac{\partial \pi_h(x_h^*(0);\tau)}{\partial \tau} = 0.$$

The derivative $\frac{\partial}{\partial x} \pi_h(x; 0) \Big|_{x=x_h^*(0)} = \frac{d}{dx} \Pr(\theta \ge \theta^* | x) \Big|_{x=x_h^*(0)}$ is positive and thus sign $\left(\frac{dx_h^*(\tau)}{d\tau}\right) = -\operatorname{sign}\left(\frac{\partial \pi_h(x_h^*(0); \tau)}{\partial \tau}\right)$. We derive a simple characterization for the right-hand side:

$$\frac{\partial \pi_h \left(x_h^*(0); \tau \right)}{\partial \tau} = \mathbf{E} \left[\left. \frac{\partial}{\partial \tau} V_{h1}(x_{t+1}; \tau) - \frac{\partial}{\partial \tau} V_{h0}(x_{t+1}; \tau) \right| x_h^*(0) \right],$$

and $\frac{\partial}{\partial \tau} V_{ha}(x_{t+1};\tau) = - \mathbb{E} \left[P_{z(ha,\mathbf{x};s(0))} | x_{t+1} \right]$ where $P_z = \sum_{t'=2}^{T} \mathbf{1}_{a(z,t')=r(z(t'-1))}$ denotes the number of punished actions on the path z. Thus

$$\frac{\partial \pi_h \left(x_h^*(0); \tau \right)}{\partial \tau} = \mathbf{E} \left[\left. P_{z(h0,\mathbf{x};s(0))} - P_{z(h1,\mathbf{x};s(0))} \right| x_h^*(0) \right].$$

This further simplifies because the best response is history-independent for $\tau = 0$; $s_h(x) = s_{h'}(x)$ when |h| = |h'|. Therefore the expected number of punished actions at rounds $t + 2, \ldots, T$ is equal across $z(h1, \mathbf{x}; s(0))$ and $z(h0, \mathbf{x}; s(0))$. Thus

$$\frac{\partial \pi_h \left(x_h^*(0); \tau \right)}{\partial \tau} = \begin{cases} 1 & \text{if } a\left(h, |h| \right) = 1, \\ -1 & \text{if } a(h, |h|) = 0, \end{cases}$$

as needed.

Next lemma characterizes how the change in the strategy affects critical volatility.

Lemma 7. For all histories h ending with action 0 or 1, $\frac{\partial}{\partial x_h^*}v^* \leq 0$ or $\frac{\partial}{\partial x_h^*}v^* \geq 0$, respectively. The inequality is strict for those h for which the inequality (8) is strict.

Proof of Lemma 7. Recall that R_h is the event that an agent reaches path h. We can write critical volatility v^* for $\tau = 0$ as

$$\sum_{h:|h|=t-1} \Pr(R_h|\theta^*) \left(\int_{-\infty}^{x_h^*} \operatorname{E}\left[v_{z(h0,\mathbf{x};s(0))} \middle| x_t \text{ and } \theta^* \right] g_h(x_t) dx_t + \int_{x_h^*}^{\infty} \operatorname{E}\left[v_{z(h1,\mathbf{x};s(0))} \middle| x_t \text{ and } \theta^* \right] g_h(x_t) dx_t \right)$$

where $g_h(x_t)$ is density of x_t conditional on θ^* and on R_h . Therefore

$$\frac{\partial}{\partial x_h^*} v^* = \Pr(R_h | \theta^*) g_h(x_h^*) \operatorname{E} \left[v_{z(h0, \mathbf{x}; s(0))} - v_{z(h1, \mathbf{x}; s(0))} \middle| x_h^* \text{ and } \theta^* \right].$$

We only need to analyze the sign of $\mathbb{E}\left[v_{z(h0,\mathbf{x};s(0))} - v_{z(h1,\mathbf{x};s(0))} | x_h^* \text{ and } \theta^*\right]$ as $\Pr(R_h|\theta^*)g_h(x_h^*)$ is positive. When $\tau = 0$, the distribution of the continuation histories h' is independent of the action played at round t:

$$\Pr(z(h0, \mathbf{x}; s(0)) = h0h' \mid x_h^* \text{ and } \theta^*) = \Pr(z(h1, \mathbf{x}; s(0)) = h1h' \mid x_h^* \text{ and } \theta^*),$$

for all $h' \in \{0, 1\}^{T-|h|-1}$ because s(0) is history-independent.

Consider h ending with action 0.¹⁸ By A3, $v_{h0h'} \leq v_{h1h'}$ for all h', and thus

$$\mathbb{E}\left[v_{z(h0,\mathbf{x};s(0))} - v_{z(h1,\mathbf{x};s(0))} \middle| x_h^* \text{ and } \theta^*\right] \le 0,$$
(26)

for any conditional distribution of the continuation histories h'. If the strict inequality holds, $v_{h0h'} < v_{h1h'}$ for some h' then (26) applies with strict inequality.

Combining the two lemmas, the effect of taxes on the critical volatility is

$$\frac{dv^*}{d\tau} = \sum_{h} \frac{\partial v^*}{\partial x_h^*} \frac{d}{d\tau} x_h^*(\tau) \le 0,$$

 $^{^{18}}$ We omit the essentially identical case of histories ending with 1.

as the summands are non-positive. If (8) is strict for some h then the summand of h is negative.

A.5 Proofs for Section 7

Proof of Proposition 8. Critical volatility $v^* = v_{\emptyset}/T$ where v_{\emptyset} is defined recursively as:¹⁹

$$v_{\emptyset} = p_{\emptyset}^* v_0 + (1 - p_{\emptyset}^*) v_1,$$

$$v_h = p_h^* (a(h, |h|) + v_{h0}) + (1 - p_h^*) (1 - a(h, |h|) + v_{h1}), \text{ when } |h| = 1 \dots T - 2,$$

$$v_h = p_h^* a(h, |h|) + (1 - p_h^*) (1 - a(h, |h|)), \text{ when } |h| = T - 1,$$

where probabilities p_h^* are given by the indifference condition (6): $p_{\emptyset}^* = \frac{c+\tau}{1+2\tau}$. For all histories of length $1, \ldots, T-2$, $p_h^* = \frac{c+2\tau}{1+2\tau}$ 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 hending with 1.

Proposition 4 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 $\frac{dv_{\emptyset}}{d\tau}\Big|_{\tau=0} < 0$. We will prove this by induction over the length of the history h.

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.

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

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 $\tau = 0$.

¹⁹Variable 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.

Finally,

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

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 $\tau = 0$.

Proof of Lemma 2. Let $(a_1, \ldots, a_K; p_1, \ldots, p_K)$ 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, \ldots, Z_K , each with mean a_k , such that L' can be identified with the compound lottery $(Z_1, \ldots, Z_K; p_1, \ldots, p_K)$. We write $L' \succ L$ if L' is a mean preserving spread of L.

We say that a set of real numbers $\{b_1, \ldots, b_K\}$ is a spread of $\{a_1, a_2\}$ if for each $k \in \{1, \ldots, K\}$, $b_k \leq a_1$ 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 8.** 1. If $\{b_1, \ldots, b_K\}$ is a spread of $\{a_1, a_2\}$, lotteries L and L' have equal means, and supports $\{a_1, a_2\}$ and $\{b_1, \ldots, b_K\}$, respectively, then $L' \succ L$.
 - 2. Suppose $L'_k \succ L_k$ for all $k = 1, \ldots, K$. Let $L = (L_1, \ldots, L_K; p_1, \ldots, p_K)$ and $L' = (L'_1, \ldots, L'_K; p_1, \ldots, p_K)$ be compound lotteries. Then $L' \succ L$.

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

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

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

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

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 $t = 0, \ldots, T$. 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, ..., 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 $(z^h_g)_{g:|g|=t-1}$, with mean $\mathbb{E}[Z_h] = S_h$, 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 $\hat{Q}_h = \left(\left(Q'_g \right)_{g:|g|=t-1}; (z_g^h)_{g:|g|=t-1} \right)$. It is constructed from Z_h by replacing each outcome S'_g by binary lottery $Q'_g = \left(S'_{g0}, S'_{g1}; p'_g, 1 - p'_g \right)$ 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

$$\left(\left(\hat{Q}_{h}\right)_{h:|h|=t-1};(l_{h})_{h:|h|=t-1}\right)$$
, and $\left(\left(Q_{h}\right)_{h:|h|=t-1};(l_{h})_{h:|h|=t-1}\right)$,

respectively.

Using the second statement of Lemma 8, $L'_t \succ L_t$ 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 8, 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 $\min_g S_g(\tau') \leq S_{h0}(\tau) < S_{h1}(\tau) \leq \max_g S_g(\tau')$.

Proof of Proposition 10. Recall that in the critical state θ^* , and in the limit $\sigma \to 0$, 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 ,

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

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, $\theta^*(T) \to 1 - \varphi(1-c) + \lambda \frac{2(1-c)c}{1+2\tau}$, as $T \to \infty$.

A.6 Proof for Subsection 8.1

Proof of Proposition 11. Let $s^* = s(u, 1)$ be the best response to θ^* under the uniform prior u and the scaling parameter $\sigma = 1$. We will prove that (rescaled) best response under a prior ϕ converges to s^* :

$$\lim_{\sigma \to 0} s_h \left(\theta^* + \sigma \varepsilon_{|h|+1}; \phi, \sigma \right) = s_h^* \left(\theta^* + \varepsilon_{|h|+1} \right), \tag{27}$$

for all histories h and any error $\varepsilon_{|h|+1}$.

The convergence of the strategies implies that, for each realization of errors, agent's actions in the critical state under prior ϕ converge to the actions under the uniform prior: denote by ε the vector ($\varepsilon_1, \ldots, \varepsilon_T$), and use **e** for the unit vector. Recall that $z(\mathbf{x}; s)$ is the terminal history reached by type **x** under strategy *s*. The convergence of the strategy, (27), implies that

$$\lim_{\sigma \to 0} z \big(\theta^* \mathbf{e} + \sigma \varepsilon; s(\phi, \sigma) \big) = z \big(\theta^* \mathbf{e} + \varepsilon; s^* \big).$$

Thus, denoting the joint errors' density by $f(\varepsilon)$,

$$\lim_{\sigma \to 0} b^*(\phi, \sigma) = \lim_{\sigma \to 0} \int b_{z\left(\theta^* \mathbf{e} + \sigma\varepsilon; s(\phi, \sigma)\right)} f(\varepsilon) d\varepsilon = \int b_{z(\theta^* \mathbf{e} + \varepsilon; s^*)} f(\varepsilon) d\varepsilon.$$

The last expression is the critical investment under the uniform prior, which equals the success premium S by Theorem 1 in the main text.

To prove (27), we will examine the belief of type x_t in round t about her next round signal x_{t+1} . Let $\tilde{f}_t(x_{t+1}, x_t; \phi, \sigma)$ denote the conditional density of x_{t+1} , given x_t , under the prior $\phi(\theta)$, and the scaling parameter σ . We will use that for all $t = 1, \ldots, T$

$$\lim_{\sigma \to 0} \tilde{f}_t \left(\theta^* + \sigma \varepsilon_{t+1}, \theta^* + \sigma \varepsilon_t; \phi, \sigma \right) = f_t (\varepsilon_t - \varepsilon_{t+1}), \tag{28}$$

recalling that f_t is the density of the instantaneous error $\eta_t = \varepsilon_t - \varepsilon_{t+1}$. Let us prove (28) for $t = 1, \ldots, T - 1$. The case for t = T is similar, only simpler. Let α_t be the density of ε_t .

$$\tilde{f}_t\left(\theta^* + \sigma\varepsilon_{t+1}, \theta^* + \sigma\varepsilon_t; \phi, \sigma\right) = \frac{\int \phi(\theta^* + \sigma\varepsilon_{t+1} - \sigma\tilde{\varepsilon}_{t+1})\alpha_{t+1}(\tilde{\varepsilon}_{t+1})f_t(\varepsilon_t - \varepsilon_{t+1})d\tilde{\varepsilon}_{t+1}}{\int\int \phi(\theta^* + \sigma\varepsilon_t - \sigma\eta_t - \sigma\tilde{\varepsilon}_{t+1})\alpha_{t+1}(\tilde{\varepsilon}_{t+1})f_t(\eta_t)d\eta_t d\tilde{\varepsilon}_{t+1}}.$$

The integrands are bounded, the prior ϕ is continuous and positive so, as $\sigma \to 0$, the righthand side converges to

$$\frac{\int \phi(\theta^*) \alpha_{t+1}(\tilde{\varepsilon}_{t+1}) f_t(\varepsilon_t - \varepsilon_{t+1}) d\tilde{\varepsilon}_{t+1}}{\int \int \phi(\theta^*) \alpha_{t+1}(\tilde{\varepsilon}_{t+1}) f_t(\eta_t) d\eta_t d\tilde{\varepsilon}_{t+1}} = f_t(\varepsilon_t - \varepsilon_{t+1}).$$

We now prove (27) by induction. Let $V_h^*(x) = V_h(x; u, 1)$ be the value function under the uniform prior u and scaling parameter $\sigma = 1$. Form an induction hypothesis that for all histories h of length t,

$$\lim_{\sigma \to 0} V_h \left(\theta^* + \sigma \varepsilon_{t+1}; \phi, \sigma \right) = V_h^* \left(\theta^* + \varepsilon_{t+1} \right),$$

and notice that the hypothesis is trivially satisfied for all the terminal histories, because $O(\theta^* + \sigma \varepsilon) = O(\theta^* + \varepsilon).$

The induction hypothesis implies that for all histories h of length t - 1,

$$\lim_{\sigma \to 0} V_h \left(\theta^* + \sigma \varepsilon_t; \phi, \sigma \right) = V_h^* \left(\theta^* + \varepsilon_t \right),$$

and

$$\lim_{\sigma \to 0} s_h \left(\theta^* + \sigma \varepsilon_t; \phi, \sigma \right) = s_h^* \left(\theta^* + \varepsilon_t \right).$$

The first limit holds because:

$$V_{h}\left(\theta^{*}+\sigma\varepsilon_{t};\phi,\sigma\right)=\max_{a}\int V_{ha}\left(\theta^{*}+\sigma\varepsilon_{t+1};\phi,\sigma\right)\tilde{f}_{t}\left(\theta^{*}+\sigma\varepsilon_{t+1},\theta^{*}+\sigma\varepsilon_{t};\phi,\sigma\right)d\varepsilon_{t+1},$$

which by the induction hypothesis and by (28) converges to

$$\max_{a} \int V_{ha}^{*} \left(\theta^{*} + \varepsilon_{t+1}\right) f_{t}(\varepsilon_{t} - \varepsilon_{t+1}) d\varepsilon_{t+1} = V_{h}^{*}(\theta^{*} + \varepsilon_{t}).$$

The argument for the second limit is identical.

References

- Abel, A., Dixit, A. K., Eberly, J., Pindyck, R. S., 1996. Options, the value of capital, and investment. The Quarterly Journal of Economics 111 (3), 753–77.
- 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.
- Arrow, K. J., 1968. Optimal capital policy with irreversible investment. In: Wolfe, J. (Ed.), Value, Capital and Growth, Papers in Honour of Sir John Hicks. Edinburgh University Press.
- Burdzy, K., Frankel, D., Pauzner, A., 2001. Fast equilibrium selection by rational players living in a changing world. Econometrica 69 (1), 163–189.

- Caballero, R. J., Engel, E. M., Haltiwanger, J., 1995. Plant-level adjustment and aggregate investment dynamics. Brookings Papers on Economic Activity 26 (2), 1–54.
- 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.
- de Mello, L., 1997. Foreign direct investment in developing countries and growth: A selective survey. The Journal of Development Studies 34 (1), 1–34.
- 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.
- Froot, K. A., Scharfstein, D., Stein, J., 1993. Risk management: Coordinating corporate investment and financing policies. Journal of Finance 48, 1629–1658.
- 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.
- Hartman, R., Hendrickson, M., 2002. Optimal partially reversible investment. Journal of Economic Dynamics and Control 26 (3), 483–508.
- 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., 2013. Reversibility in dynamic coordination problems. Games and Economic Behavior 77 (1), 298–320.
- Lensink, R., Morrissey, O., 2006. Foreign direct investment: Flows, volatility, and the impact on growth. Review of International Economics 14 (3), 478–493.
- Loeper, A., Steiner, J., Stewart, C., 2012. Influential opinion leaders, unpublished paper, CERGE-EI.
- Minton, B. A., Schrand, C., 1999. The impact of cash flow volatility on discretionary investment and the costs of debt and equity financing. Journal of Financial Economics 54 (3), 423–460.
- Mobarak, A. M., 2005. Determinants of volatility and implications for economic development. Review of Economics and Statistics 87 (2), 348–361.
- 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.
- Oliva, M.-A., Rivera-Batiz, L. A., 2002. Political institutions, capital flows, and developing country growth: An empirical investigation. Review of Development Economics 6 (2), 248–62.
- Pallage, S., Robe, M. A., 2003. On the welfare cost of economic fluctuations in developing countries. International Economic Review 44 (2), 677–698.
- Ramey, G., Ramey, V. A., 1995. Cross-country evidence on the link between volatility and growth. American Economic Review 85 (5), 1138–51.
- 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.