14.462 Lecture Notes
Morris-Shin and Global Games

George-Marios Angeletos
MIT Department of Economics

Spring 2004

1 Framework

Actions, Outcomes and Payoffs. There are two possible regimes, the status quo and an alternative. There is a measure-one continuum of agents, indexed by $i \in [0,1]$. Each agent can choose between an action that is favorable to the alternative regime and an action is favorable to the status quo. We call these actions, respectively, "attack" and "not attack". All agents move simultaneously.

We denote the regime outcome with $R \in \{0,1\}$, where $R = 0$ represents survival of the status quo and $R = 1$ represents collapse. We similarly denote the action of an agent with $a_i \in \{0,1\}$, where $a_i = 0$ represents "not attack" and $a_i = 1$ represents "attack".

The payoff from not attacking is normalized to zero. The payoff from attacking is $b > 0$ if the status quo is abandoned and $-c < 0$ otherwise. Hence, the utility of agent i is

$$u_i = a_i(bR - c).$$

Finally, the status quo is abandoned ($R = 1$) if and only if

$$A \geq \theta,$$

where $A \equiv \int a_i di \in [0,1]$ denotes the mass of agents attacking and $\theta \in \mathbb{R}$ parameterizes the exogenous strength of the status quo (or the quality of the economic fundamentals). Let $\underline{\theta} \equiv 0$ and $\overline{\theta} \equiv 1$.

Complementarity. Note that the actions of the agents are strategic complements, since it pays for an individual to attack if and only if the status quo collapses and, in turn, the status quo collapses if and only if a sufficiently large fraction of the agents attacks. This coordination problem is the heart of the model. To see this more clearly, rewrite the payoff of agent i as

$$u_i = U(a_i, A, \theta) = \begin{cases} a_i(b-c) & \text{if } A \geq \theta \\ -a_i c & \text{if } A < \theta \end{cases}$$

Assuming θ is known, the best response is

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

Note that both U and g are increasing in A, so that we have both a positive externality and a complementarity. In particular, the complementarity is of the "strong" form we discussed in the first lecture.

Interpretations. This simple model can capture the role of coordination and multiplicity of equilibria in a variety of interesting applications. For instance, in models of self-fulfilling currency crises (Obstfeld, 1986, 1996; Morris and Shin, 1998), there is a central bank interested in maintaining a currency peg and a large number of speculators, with finite wealth, deciding whether to attack the currency or not. In this context, a "regime change" occurs when a sufficiently large mass of speculators attacks the currency, forcing the central bank to abandon the peg.

In models of self-fulling bank runs, a "regime change" occurs once a sufficiently large number of depositors decide to withdraw their deposits, relative to liquid resources available to the system, forcing the bank to suspend its payments. Similarly, in models of self-fulfilling debt crises (Calvo, 1988; Cole and Kehoe, 1996; Morris and Shin; 2003), a "regime change" occurs when a lender fails to obtain refinancing by a sufficiently large fraction of its creditors.

Finally, Atkeson (2000) interprets the model as describing riots. The potential rioters may or may not overwhelm the police force in charge of containing social unrest depending on the number of the rioters and the strength of the police force.

Information. Suppose for a moment that θ were commonly known by all agents. For $\theta \leq \underline{\theta}$, the fundamentals are so weak that the regime is doomed with certainty and the unique equilibrium is every agent attacking. For $\theta > \overline{\theta}$, the fundamentals are so strong that the regime can survive an attack of any size and the unique equilibrium is every agent not attacking.

For intermediate values, $\theta \in (\underline{\theta}, \overline{\theta}]$, the regime is sound but vulnerable to a sufficiently large attack and there are multiple equilibria sustained by self-fulfilling expectations. In one equilibrium, individuals expect everyone else to attack, they then find it individually optimal to attack, the status quo is abandoned and expectations are vindicated. In another, individuals expect no one else to attack, they then find it individually optimal not to attack, the status quo is spared and expectations are again fulfilled. The interval $(\underline{\theta}, \overline{\theta}]$ thus represents the set of "critical fundamentals" for which multiple equilibria are possible under common knowledge.

Implicitly, each equilibrium is sustained by different self-fulfilling expectations about what other agents do. With common knowledge, in equilibrium individuals can perfectly forecast each other actions and coordinate on multiple courses of action. Following Morris and Shin (1998), we assume that θ is never common knowledge and that individuals instead have private noisy information about θ. Private information serves as an anchor for individual's actions that may avoid the indeterminacy of expectations about others actions.

Initially agents have a common prior about θ; for simplicity, we let this prior be (degenerate) uniform over the entire real line. Agent i then observes a private signal

$$x_i = \theta + \xi_i,$$

where the idiosyncratic noise ξ_i is $\mathcal{N}(0, \sigma_x^2)$ with $\sigma_x > 0$ and is independent of θ. The signal x_i is thus a sufficient statistic for the private information of an agent.

Note that because there is a continuum of agents the information contained by the entire economy, $(x_i)_{i \in [0,1]}$, is enough to infer the fundamental θ. However, this information is dispersed throughout the population, which is the key feature of the Morris-Shin framework.

Finally, agents may also have access to some public information. In particular, agents observe an *exogenous* public signal $z = \theta + v$, where $v \sim \mathcal{N}(0, \sigma_z^2)$. The public noise v is distributed independently of θ and the private noise ξ.

Equilibrium. An economy is parametrized by the standard deviations (σ_x, σ_z), or equivalently the precisions (α_x, α_z), where $\alpha_x = \sigma_x^{-2}$ and $\alpha_z = \sigma_z^{-2}$. A symmetric Bayesian equilibrium is defined as follows

Definition 1 *An equilibrium is a strategy $a(.)$ and an aggregate attack $A(.)$ such that:*

$$\begin{aligned} a(x,z) &\in \arg\max_a \mathbb{E}\left[U(a, A(\theta,z), \theta) \mid x, z\right] \\ A(\theta,z) &= \int a(x,z)\sqrt{\alpha_x}\phi\left(\sqrt{\alpha_x}[x-\theta]\right) dx. \end{aligned}$$

2 Monotone Equilibria

We start by considering *monotone or threshold equilibria*, that is, equilibria in which $a(x,z)$ is monotonic in x.

In a monotone equilibrium, for any realization of z, there is a threshold $x^*(z)$ such that an agent attacks if and only if $x \leq x^*(z)$. By implication, the aggregate size of the attack is decreasing in θ, so that there is also a threshold $\theta^*(z)$ such that the status quo is abandoned if and only if $\theta \leq \theta^*(z)$. A monotone equilibrium is identified by x^* and θ^*. In step 1, below, we characterize the equilibrium θ^* for given x^*. In step 2, we characterize the equilibrium x^* for given θ^*. In step 3, we characterize both conditions and examine equilibrium existence and uniqueness.

Step 1. For given realizations of θ and z, the aggregate size of the attack is given by the mass of agents who receive signals $x \leq x^*(z)$. That is,

$$A(\theta, z) = \Phi\left(\sqrt{\alpha_x}(x^*(z) - \theta)\right),$$

where $\alpha_x = \sigma_x^{-2}$ is the precision of private information. Note that $A(\theta, z)$ is decreasing in θ, so that regime change occurs if and only if $\theta \leq \theta^*(z)$, where $\theta^*(z)$ is the unique solution to

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

Rearranging we obtain:

$$x^*(z) = \theta^*(z) + \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}(\theta^*(z)). \tag{1}$$

Step 2. Given that regime change occurs if and only if $\theta \leq \theta^*(z)$, the payoff of an agent is

$$\mathbb{E}\left[\, U(a, A(\theta, z), \theta) \mid x, z \,\right] = a(b \Pr\left[\theta \leq \theta^*(z) \mid x, z\right] - c).$$

Let $\alpha_x = \sigma_x^{-2}$ and $\alpha_z = \sigma_z^{-2}$ denote, respectively, the precision of private and public information. The posterior of the agent is

$$\theta \mid x, z \;\sim \mathcal{N}\left(\; \delta x + (1-\delta)z \;,\; \alpha^{-1} \;\right),$$

where $\delta \equiv \alpha_x/(\alpha_x + \alpha_z)$ is the relative precision of private information and $\alpha \equiv \alpha_x + \alpha_z$ is the overall precision of information. Hence, the posterior probability of regime change is

$$\Pr\left[\, \theta \leq \theta^*(z) \mid x, z \,\right] = 1 - \Phi\left(\sqrt{\alpha}(\delta x + (1-\delta)z - \theta^*(z))\right),$$

which is monotonic in x. It follows that the agent attacks if and only if $x \leq x^*(z)$, where $x^*(z)$ solves the indifference condition

$$b \Pr\left[\, \theta \leq \theta^*(z) \mid x^*(z), z \,\right] = c.$$

Substituting the expression for the posterior and the definition of δ and α, we obtain:

$$\Phi\left(\sqrt{\alpha_x+\alpha_z}\left(\frac{\alpha_x}{\alpha_x+\alpha_z}x^*(z)+\frac{\alpha_z}{\alpha_x+\alpha_z}z-\theta^*(z)\right)\right)=\frac{b-c}{b}. \tag{2}$$

Step 3. Combining (1) and (2), we conclude that $\theta^*(z)$ can be sustained in equilibrium if and only if it solves

$$G\left(\theta^*(z),z\right)=g, \tag{3}$$

where $g=\sqrt{1+\alpha_z/\alpha_x}\Phi^{-1}\left(1-c/b\right)$ and

$$G(\theta,z)\equiv\frac{\alpha_z}{\sqrt{\alpha_x}}\left(z-\theta\right)+\Phi^{-1}\left(\theta\right).$$

With $\theta^*(z)$ given by (3), $x^*(z)$ is then given by (1). We are now in a position to establish existence and determinacy of the equilibrium by considering the properties of the function G. Note that, for every $z\in\mathbb{R}$, $G(\theta,z)$ is continuous in θ, with $G(\underline{\theta},z)=-\infty$ and $G(\overline{\theta},z)=\infty$, which implies that there necessarily exists a solution and any solution satisfies $\theta^*(z)\in(\underline{\theta},\overline{\theta})$. This establishes existence; we now turn to uniqueness. Note that

$$\frac{\partial G(\theta,z)}{\partial\theta}=\frac{1}{\phi(\Phi^{-1}\left(\theta\right))}-\frac{\alpha_z}{\sqrt{\alpha_x}}$$

Since $\max_{w\in\mathbb{R}}\phi(w)=1/\sqrt{2\pi}$ then if $\alpha_z/\sqrt{\alpha_x}\leq\sqrt{2\pi}$ we have that G is strictly increasing in θ, which implies a unique solution to (3). If instead $\alpha_z/\sqrt{\alpha_x}>\sqrt{2\pi}$, then G is non-monotonic in θ and there is an interval $(\underline{z},\overline{z})$ such that (1) admits multiple solutions $\theta^*(z)$ whenever $z\in(\underline{z},\overline{z})$ and a unique solution otherwise. We conclude that monotone equilibrium is unique if and only if $\alpha_z/\sqrt{\alpha_x}\leq\sqrt{2\pi}$.

We summarize these results in the following proposition.

Proposition 2 (Morris-Shin) *Let σ_x and σ_z denote the standard deviations of the private and the public noise, respectively. There always exists a monotone equilibrium and it is unique if and only if $\sigma_x/\sigma_z^2\leq\sqrt{2\pi}$.*

Finally, consider the limits as $\sigma_x \to 0$ for given σ_z, or $\sigma_z \to \infty$ for given σ_x. In either case, $\alpha_z/\sqrt{\alpha_x} \to 0$ and $\sqrt{(\alpha_x + \alpha_z)/\alpha_x} \to 1$. Condition (3) then implies that $\theta^*(z) \to \widehat{\theta} = 1 - c/b$, for all z. This proves the following result, which we refer to as the *Morris-Shin limit result*:

Proposition 3 (Morris-Shin limit) *In the limit as either $\sigma_x \to 0$ for given σ_z, or $\sigma_z \to \infty$ for given σ_x, there is a unique monotone equilibrium in which the regime changes if and only if $\theta \leq \widehat{\theta}$, where $\widehat{\theta} = 1 - c/b \in (\underline{\theta}, \overline{\theta})$.*

3 Iterated Dominance Argument

The results above established that there exists a unique monotone equilibrium whenever the noise in private information is small enough. These results, however, left open the possibility that there are other non-monotone equilibria. We now prove the much stronger result that there is no other equilibrium and, what is more, that the equilibrium is dominance solvable.

To simplify the exposition, consider the case that there is no public information: $\sigma_z = \infty$ ($\alpha_z = 0$), implying $\delta = 1$ and $\alpha = \alpha_x$. We can thus drop z and denote the strategy by $a(x)$ and the aggregate attack by $A(\theta)$.

For any $\widehat{x} \in [-\infty, +\infty]$, let $A_{\widehat{x}}(\theta)$ denote the size of aggregate attack when (almost every) agent attacks if and only if $x \leq \widehat{x}$.Next, define the function

$$V(x, \widehat{x}) = \mathbb{E}[\, U(1, A_{\widehat{x}}(\theta), \theta) - U(0, A_{\widehat{x}}(\theta), \theta) \mid x \,].$$

This is the utility difference between attacking and not attacking for an agent who has private information x and expects the other agents to attack if and only if their signals fall below $\widehat{x}$. In our model,

$$A_{\widehat{x}}(\theta) = \Phi\left(\sqrt{\alpha_x}\,[\widehat{x} - \theta]\right)$$

and

$$V(x, \widehat{x}) = b - b\Phi\left(\sqrt{\alpha_x}[x - \widehat{\theta}]\right) - c,$$

where $\widehat{\theta} = \widehat{\theta}(\widehat{x})$ is the unique solution to $A_{\widehat{x}}(\widehat{\theta}) = \widehat{\theta}$, or equivalently the inverse of

$$\widehat{x} = \widehat{\theta} + \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}(\widehat{\theta}).$$

Note $\widehat{\theta}$ is increasing in $\widehat{x}$. It follows that $V(x, \widehat{x})$ is increasing in $\widehat{x}$: The more aggressive the other agents are, the higher the expected payoff from attacking. Moreover, $V(x, \widehat{x})$ is decreasing in x : The higher the private signal, the lower the expected payoff from attacking.

Next, note that $V(x, \widehat{x})$ is continuous in x and satisfies $V(x, \widehat{x}) \to b - c > 0$ as $x \to -\infty$ and $V(x, \widehat{x}) \to -c < 0$ as $x \to +\infty$. We can thus define a function h such that $x = h(\widehat{x})$ is the unique solution to $V(x, \widehat{x}) = 0$ with respect to x. The interpretation of $h(\widehat{x})$ is simple: When agents $j \neq i$ attack if and only if $x_j \leq \widehat{x}$, agent i finds it optimal to attack if and only if $x_i \leq h(\widehat{x})$. Because $V(x, \widehat{x})$ is continuous in both arguments, decreasing in x, and increasing in $\widehat{x}$, the function $h(\widehat{x})$ is continuous and increasing in $\widehat{x}$. Finally, note that h has a unique fixed point $x^* = h(x^*)$ and this fixed point is indeed the threshold of the unique monotone equilibrium we constructed in the previous section.

Construct a sequence $\{\underline{x}_k\}_{k=0}^{\infty}$ by $\underline{x}_0 = -\infty$ and $\underline{x}_k = h(\underline{x}_{k-1})$ for all $k \geq 1$. In particular, letting $\underline{\theta}_{k-1}$ be the solution to

$$\underline{x}_{k-1} = \underline{\theta}_{k-1} + \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}(\underline{\theta}_{k-1}),$$

we have

$$V(x, \underline{x}_{k-1}) = b - b\Phi\left(\sqrt{\alpha_x}[x - \underline{\theta}_{k-1}]\right) - c$$

and thus

$$\underline{x}_k = \underline{\theta}_{k-1} + \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}\left(\frac{b-c}{b}\right).$$

Hence, $\underline{x}_0 = -\infty$, $\underline{\theta}_0 = 0$, $\underline{x}_1 = \frac{1}{\sqrt{\alpha_x}}\Phi^{-1}\left(\frac{b-c}{b}\right)$, and so on. Clearly, the sequence $\{\underline{x}_k\}_{k=0}^{\infty}$ is increasing and bounded above by x^*. Hence, the sequence $\{\underline{x}_k\}_{k=0}^{\infty}$ converges to some $\underline{x}$. By continuity of h, the limit $\underline{x}$ must be a fixed point of h. But we already proved that h has a unique fixed point. Hence, $\underline{x} = x^*$.

Next, construct a sequence $\{\overline{x}_k\}_{k=0}^{\infty}$ by $\overline{x}_0 = -\infty$ and $\overline{x}_k = h(\overline{x}_{k-1})$ for all $k \geq 1$. Note that this sequence is decreasing and bounded below by x^*. Hence, the sequence $\{\overline{x}_k\}_{k=0}^{\infty}$ converges to some $\overline{x}$. By continuity of h, the limit $\overline{x}$ must be a fixed point of h. But we already proved that h has a unique fixed point. Hence, $\overline{x} = x^*$.

But, what is the meaning of the sequences we constructed? Consider $\underline{x}_1$. If nobody else attacks, the agent finds it optimal to attack if and only if $x \leq \underline{x}_1$. By complementarity, if some people attack, the agent finds it optimal to attack *at least* for $x \leq \underline{x}_1$. That is, for $x \leq \underline{x}_1$, attacking is *dominant*. Next, consider $\underline{x}_2$. When other agents attack if and only if it is dominant for them to do so, that is, if and only if $x \leq \underline{x}_1$, then it is optimal to attack if and only if $x \leq \underline{x}_2$. By complementarity, if other agents attack at least for $x \leq \underline{x}_1$, then it is optimal to attack *at least* for $x \leq \underline{x}_2$. That is, for $x \leq \underline{x}_2$, attacking becomes dominant after the first round of deletion of dominated strategies. Similarly, for $x \leq \underline{x}_{k-1}$, attacking becomes dominant after the $k-1$ round of deletion of dominated strategies. Hence, $\{\underline{x}_k\}_{k=0}^{\infty}$ represents iterated deletion of dominated strategies "from below". Similarly, $\{\overline{x}_k\}_{k=0}^{\infty}$ represents iterated deletion of dominated strategies "from above".

To recap, the only strategies that have survived k rounds of iterated deletion of dominated strategies are functions a such that $a(x) = 1$ for all $x \leq \underline{x}_k$ and $a(x) = 0$ for $x > \overline{x}_k$; for $x \in (\underline{x}_k, \overline{x}_k)$, the value of $a(x)$ is still "free" at the k-th round. But we proved that both $\underline{x}_k$ and $\overline{x}_k$ converge to x^* as $k \to \infty$. Hence, in the limit, the only strategy that survives is the function a such that $a(x) = 1$ for $x \leq x^*$ and $a(x) = 0$ for $x > x^*$, which is precisely the strategy in the monotone equilibrium.

As long as $\alpha_z/\sqrt{\alpha_x} \leq \sqrt{2\pi}$, the above result extends to the presence of public information. (Do it as an exercise.) We conclude:

Proposition 4 *If $\sigma_x/\sigma_z^2 \leq \sqrt{2\pi}$, there is a unique equilibrium. This equilibrium is the monotone equilibrium described before and it is solvable by iterated deletion of dominated strategies.*

When instead $\alpha_z/\sqrt{\alpha_x} > \sqrt{2\pi}$, the argument breaks in the following respect. For some realizations of the public signal, the equation $V(x,x) = 0$ has three solutions (equivalently, h has three fixed points); this is what we showed in the previous section. The sequences $\{\underline{x}_k\}_{k=0}^{\infty}$ and $\{\overline{x}_k\}_{k=0}^{\infty}$ are still well defined as above. But now their limits $\underline{x}$ and $\overline{x}$ do not coincide. Instead, $\underline{x}$ is the *lowest* of the three solutions to $V(x,x) = 0$ and $\overline{x}$ is the *highest* of the three. Hence, the two monotone equilibria that we constructed in the previous section represent the least and most aggressive equilibria of the game. It is unclear whether there are also other intermediate, non-monotone equilibria.