

THE EVOLUTION OF FOCAL POINTS*

Ken Binmore
Department of Economics
University College London
Gower Street
London WC1E 6BT England

Larry Samuelson
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706-1321 USA

February 11, 2002

Abstract: Theories of focal points typically assume that games have inherent labelings or “frames” and then construct models of how players perceive and exploit these frames to identify focal equilibria. This paper asks instead how evolutionary considerations determine which aspects of a frame are likely to be monitored by the players. Efficient monitoring turns out to be an unlikely outcome.

Contents

1	Introduction	1
2	Evolution and Focal Points	2
3	A Model of Focal Points	6
4	Evolutionarily Stable Monitoring	11
4.1	Efficient equilibria	11
4.2	Invasion barriers	12
4.3	Risk Dominance	16
4.4	Stochastic Stability	17
5	Discussion	19
6	Appendix	21
6.1	One Independent Property	21
6.2	Many Independent Properties	23
6.3	The Benefit Function	24
6.4	Nonconcave $A(\alpha) - C(\alpha)$	27
6.5	Semi-Differentiability	29
6.6	Efficient Monitoring	29

*Financial support from the Leverhulme Foundation, the Economic and Social Research Council and the National Science Foundation is gratefully acknowledged.

THE EVOLUTION OF FOCAL POINTS

by Ken Binmore and Larry Samuelson

1 Introduction

Thomas Schelling once observed that focal points have done more for the theory of games than game theory has done for the theory of focal points.¹ This paper seeks to make amends by showing that evolutionary game theory can explain how contextual cues become focal in pure coordination games.

Schelling's [16, p. 55] best-known example concerns two strangers who are unable to communicate but urgently need to meet in New York City, having fixed a time but not a location for their meeting. Most of the participants in Schelling's informal experiments in the fifties chose Grand Central Station as their attempted meeting point. Of the many possible equilibrium locations, Grand Central Station was focal.

	1	2	3
1	1	0	0
2	0	1	0
3	0	0	1

	□	□	
□	1	0	0
□	0	1	0
	0	0	1

(a) (b)

Figure 1: The 3×3 game of pure coordination.

Game theorists model such a problem as a game of pure coordination. Figure 1(a) is an example with three strategies. In our analysis, we assume the players know nothing whatever about the labeling of the strategies unless they receive some additional framing information. Their ignorance extends even to such questions as which strategy corresponds to the top row or middle column, conventions that arise only in the course of constructing the game theorist's model. Whatever the players do in such a state of ignorance, the result is equivalent to their using each pure strategy with probability $\frac{1}{3}$.² They then get their worst possible equilibrium payoff.

¹At the Arne Ryde Conference held in his honor in Lund, Sweden on August 23, 1997.

²This is why the equilibrium selection theory of Harsanyi and Selten [8] selects this mixed equilibrium as the solution of a pure coordination game.

Real-life games come with *frames* that relate strategies to the context in which the game is played. Bacharach [1], Bacharach and Stahl [2], Casajus [5], Sugden [18] and others have asked how players perceive and use such a framing to locate focal points. A major concern is the extent to which a frame allows the players to coordinate on an efficient equilibrium of the game. For example, when the strategies were framed as locations in New York City, most participants in Schelling’s experiments obtained the efficient equilibrium payoff of one by choosing Grand Central Station.

How do frames arise? Why do we choose to pay close attention to some environmental signals and not others? Is framing information likely to be efficiently used? We explore an evolutionary approach to such questions, on the understanding that cultural evolution is at least as important as biological evolution in the case of human coordination. We find that efficient focal mechanisms are unlikely to survive, even when all the obvious obstacles to efficient coordination are absent. We defend this conclusion using a number of evolutionary stability criteria.

2 Evolution and Focal Points

This section presents a simple example illustrating the efficient use of framing information and the attendant evolutionary considerations.

The Square Game. Suppose that each time Nature calls for the 3×3 pure coordination game of Figure 1(a) to be played, she frames it by labeling some strategies with a square. If each strategy is independently assigned a square (or not) with probability $\frac{1}{2}$, the probability of the configuration shown in Figure 1(b) is $\frac{1}{8}$. In this configuration, the players make the best use of their framing information by coordinating on the strategy without a square, so that each receives the maximal payoff of one.

Configurations with a similar odd-man-out occur with a total probability of $\frac{3}{4}$. However, with probability $\frac{1}{4}$, either all or none of the strategies are assigned a square. The strategies are then indistinguishable, and anything the players may do is equivalent to randomizing their choice, for an expected payoff of $\frac{1}{3}$. Their overall expected payoff is therefore

$$\frac{3}{4} \times 1 + \frac{1}{4} \times \frac{1}{3} = \frac{20}{24}.$$

The Red Square Game. Suppose that Nature now enriches the frame by independently coloring some strategies red in the same way that she

assigns some strategies a square. With probability $\frac{1}{16}$, all three strategies will still be indistinguishable, so that the players can achieve a payoff of only $\frac{1}{3}$. With probability $\frac{15}{16}$, at least one strategy bears a joint property shared by no other strategy. It may be, for example, that only one strategy is red but not square. The players can then achieve perfect coordination, provided that they are party to a convention that allows them to resolve ambiguities when there is more than one such odd-man-out. Adding color to the frame therefore improves the players' overall expected payoff to

$$\frac{15}{16} \times 1 + \frac{1}{16} \times \frac{1}{3} = \frac{23}{24}.$$

The previous literature on framing in a game-theoretic context has largely concentrated on how we manage to adjudicate between rival focal points. For example, is a strategy that is red but not square more focal than a strategy that is square but not red? We bypass this important psychological issue by always assuming that the players use the framing information efficiently. For example, the information provided by the Red Square frame is used efficiently when both players employ the Shape-then-Color convention, in which a player chooses an odd-man-out in shape if there is one, and otherwise chooses an odd-man-out in color if one exists (failing which the choice is necessarily random). We confine attention to such efficient conventions in order to concentrate on an unstudied source of inefficiency—that which arises when the players pay insufficient attention to the information provided by the frame.

Costly Monitoring. Although we are seldom conscious of paying the costs, it must nevertheless be costly to monitor framing information. We are deluged with too much information to make use of it all, and the information in a framing can be processed only at the expense of ignoring something else. We capture such costs in this example by supposing that, before being confronted with the Red Square Game, players are programmed by an evolutionary process to *monitor* shape or color. Monitoring only one property costs nothing. Monitoring both incurs a cost of $c > 0$.

Choosing what to monitor is itself a strategy in an impure coordination game. The strategies in this “monitoring game” are to be understood as monitoring *both* properties, or monitoring only *one* property. In the latter case, we assume that a convention is in place ensuring that both players choose the same property. We can then represent the game as in Figure 2.

The payoff B is the expected utility received by both players in the underlying pure coordination game when a player who monitors both properties meets a player who monitors only one property. Observe that $B \leq \frac{20}{24}$,

		<i>both</i>	<i>one</i>
<i>both</i>		$\frac{23}{24} - c$	B
	$\frac{23}{24} - c$	$B - c$	$B - c$
<i>one</i>		B	$\frac{20}{24}$
	B	$\frac{20}{24}$	$\frac{20}{24}$

Figure 2: A monitoring game.

because a player cannot coordinate more with an opponent than a player who monitors the same information as the opponent. In our example, we assume that

$$B = \frac{20}{24},$$

which is achieved when a player who monitors both properties uses the Shape-then-Color convention, and a player who monitors only one property uses the Shape-only convention.³ In more complex games, this upper bound on B cannot always be attained.

When $c > \frac{1}{8}$, the cost of monitoring a second property outweighs the value of the additional information and it is a strictly dominant strategy to monitor only one property. If $c < \frac{1}{8}$, each strategy corresponds to a strict Nash equilibrium. The efficient equilibrium in this case requires monitoring both properties, but we shall see that evolutionary arguments may favor the inefficient equilibrium in which only one property is monitored.

Evolutionary Dynamics. In redirecting attention from games of pure coordination to the monitoring games within which they are embedded, we replace one equilibrium selection problem by another. So what progress has been made? Our approach is to regard behavior in the monitoring game as the outcome of an evolutionary process rather than of conscious choice. This allows us to bring evolutionary considerations to bear on the equilibrium selection problem.

When $c < \frac{1}{8}$, the monitoring game of Figure 2 has three Nash equilibria. The two pure equilibria correspond to the strategies *both* and *one*. Which will evolution select? Both are *strict* Nash equilibria, and so both are evolutionarily stable (Maynard Smith and Price [11], Selten [17]). It is therefore necessary to look more deeply into the dynamics of the evolutionary process that motivates the definition of evolutionary stability.

³If the latter player uses the Color-only convention, then $B = \frac{11}{24} < \frac{20}{24}$.




Figure 3: Basins of attraction.

If the stochastic element of an evolutionary process is sufficiently small, then a deterministic model such as the replicator dynamics of evolutionary biology will suffice to predict with high probability the equilibrium of a game that the dynamic system will approach in the first instance (Binmore, Samuelson and Vaughan [3]). The mixed equilibrium in our example is shown as the point M in the phase diagram of Figure 3. It corresponds to a distribution of the underlying population in which a fraction $8c$ of the population play *both*, and the remainder play *one*. The arrows in the phase diagram show the evolutionary flows of any monotonic deterministic evolutionary dynamics (Samuelson and Zhang [15]). The population states in which everybody plays *both* or everybody plays *one* are stable with respect to such dynamics, but M is unstable. In the long run, the system converges on one of the two stable equilibria, but the equilibrium selected depends on where the system starts.

If we are interested in very long periods of time, then even very small stochastic perturbations must be explicitly considered, since they will occasionally bounce the system from one equilibrium's basin of attraction into the other (Young [20], Kandori, Mailath and Rob [10]). We will therefore find the system sometimes near *both*, and sometimes near *one*. We say that an equilibrium is selected in the ultralong run if it is almost certain that the system will be found in its vicinity when the perturbations are made sufficiently small.⁴

The equilibrium with the larger basin of attraction in our example is said to be *risk dominant* (Harsanyi and Selten [8]). Whichever equilibrium assigns the players the larger payoff is *payoff dominant*. In our example, *both* is payoff dominant, and *one* is risk dominant when $\frac{1}{16} < c < \frac{1}{8}$. Since evolutionary dynamics typically select the risk-dominant equilibrium in the

⁴The expected waiting times before ultralong-run predictions are realized may be very long, even by geological standards. However, Binmore, Samuelson and Vaughan [3] show that more realistic models based on the kind of reasoning that leads to the replicator dynamics have expected waiting times that make predictive sense. Young [21] has also shown that expected waiting times can be greatly reduced when the process that selects players to play a game is viscous, which means that it favors matches between neighbors.

ultralong run,⁵ we should then expect to see only one property being monitored, although it is efficient to monitor two properties.

The cause of this inefficiency is not that there is a mismatch between individual and joint incentives—the efficient outcome is an equilibrium. It lies in an inefficient balancing of the costs and benefits of monitoring information.

3 A Model of Focal Points

The rest of the paper shows that the evolutionary selection of inefficient equilibria in monitoring games is a general phenomenon. This section outlines our model.

Pure Coordination. Our techniques apply to multiplayer games of impure coordination, but we sharpen the results by considering only two-player games of *pure coordination*. Each player has the strategy set $\{1, 2, \dots, N\}$ and the payoff function

$$\pi(n_1, n_2) = \begin{cases} 1 & \text{if } n_1 = n_2 \\ 0 & \text{otherwise} \end{cases},$$

where n_1 and n_2 are the choices of players 1 and 2. The strategy labels $\{1, 2, \dots, N\}$ cannot themselves be used to coordinate behavior (Crawford and Heller [6]). In keeping with Schelling’s New York example, we are particularly interested in the case when N is large.

Frames. We postulate that each pure strategy of the game potentially satisfies each of a number of *properties*: P_1, P_2, \dots, P_M . Examples of possible properties include *square, first, red, central, brightest, top-left*.

Formally, a property is a random function $P_i : \{1, \dots, N\} \rightarrow [0, 1]$, where $P_i(n) = 1$ means that strategy n has property i . Properties are assigned randomly in our model to reflect the fact that players are likely to encounter the coordination game in many different contexts. A *frame* for the game is the random process that determines which properties are

⁵Young [20] and Kandori, Rob and Mailath [10] study stochastic dynamics in which the risk-dominant equilibrium is always selected. Binmore, Samuelson and Vaughan [3] examine a model in which *both* continues to be selected, despite being risk dominated, when c is slightly larger than $\frac{1}{16}$, but in which *one* is (inefficiently) selected for larger values of c . Some evolutionary models always select the payoff-dominant equilibrium (Robson and Vega Redondo [12]), but these require rather special circumstances.

realized each time that Nature decides that the coordination game is to be played. Although properties are assigned independently across strategies in our examples, we make no general independence assumptions concerning either the realization of a given property across strategies or the realizations of different properties. Otherwise we would not be able to accommodate properties like *brightest* or *top-left*.

Nature is given her best chance at creating efficient focal points by assuming that players who pay attention to the same framing information use it with maximal efficiency. They first use the framing information to identify sets of pure strategies in the underlying game of pure coordination that are of minimal size. A commonly held convention then selects one of these minimal sets. Each pure strategy in the focal set is then played with equal probability. We say that the selected set is *focal*.

One independent property. Suppose the frame independently assigns a single property to each strategy with probability p . If players costlessly monitor only one property, what value of p maximizes their expected payoff?

Section 6.1 shows that the optimal p for a coordination game with N strategies is either $p^*(N)$ or $1 - p^*(N)$, where $p^* = \frac{1}{2}N$ when $N \leq 4$, but

$$p^*(N) \approx 1.3 N^{-1}$$

when N is sufficiently large. This result for large N illustrates the balance that a useful coordination property must strike. A property that is likely to be shared by too many strategies is unhelpful, because it is too ambiguous, frequently forcing the players to mix over large sets of strategies satisfying the property. A property that is sufficiently rare as to be unlikely to be shared may often hold for no strategy at all, providing no help with coordination. For large N , the optimal property balances these forces by setting the expected number of realizations to be just over one.

Many independent properties. Suppose that there are M independently assigned properties, with any given strategy satisfying property P_i with probability p_i . Section 6.2 shows that if the probabilities p_i can be chosen at will, then perfectly efficient coordination is approached very rapidly as the number of properties grows. With M properties, the expected payoff for large values of N is at least

$$1 - \left(1 - \frac{1}{e}\right)^{2^{M-1}}.$$

With five properties, for example, the efficiency loss is less than 0.001, no matter how large N may be.

This nearly perfect coordination is achieved by an optimal frame consisting of $M - 1$ properties with $p_i = \frac{1}{2}$, and one property with either p_M or $1 - p_M$ equal to $p^*(N)2^{M-1}$, where $p^*(N)$ is the same as in the case of only one property. The $M - 1$ properties with $p_i = \frac{1}{2}$ are very likely to partition the N strategies into 2^{M-1} subsets of approximately the same size. The final property can then be viewed as a chance at isolating a single strategy in one of these subsets to serve as a focal point. The value $p^*(N)2^{M-1}$ allows this chance to be taken optimally in each of the 2^{M-1} subsets of size $N/2^{M-1}$. Except when M is quite small, the risk of these 2^{M-1} independent attempts failing to isolate any strategy in any of these subsets is tiny.

Paying Attention The finding that a handful of well-chosen independent properties can lead to nearly perfect coordination contrasts with the observation that focal points are often difficult to locate. Grand Central Station may have seemed an obvious meeting point for New Yorkers in the fifties, but where is the corresponding focal point in contemporary Columbus?

One difficulty is that we seldom have the luxury of working with optimally designed properties. Nor is it costless to monitor the information that a frame provides. For example, when a frame labels strategies with words, we do not ask ourselves whether one of these might be an anagram of a Brazilian poet's name, presumably because the added benefit of monitoring such properties closely does not justify the cost. We therefore introduce monitoring costs that increase with the amount of attention that a player pays to various properties.

Costs and Benefits. The *intensity* $\alpha_i \in \mathbb{R}_+$ with which a player monitors each property P_i is chosen by evolution. The intensities $\alpha_1, \dots, \alpha_M$ determine the probabilities with which a player recognizes whatever realizations (if any) of each property have occurred in a framing of the coordination game.⁶ A monitoring *strategy* is a vector α of intensities determined before anything else happens.

Just as the realizations of properties need not be independent, the event that two properties are observed need not be independent. A player who observes which strategy is *largest* (if any) may also be likely to observe

⁶Hence, monitoring a property with positive intensity does not ensure that some strategy exhibits the property, nor that the property will be observed even when it holds for some strategies.

smallest.

A player using the monitoring strategy α who faces a player using the monitoring strategy β derives a *benefit* $B(\alpha, \beta)$, which is the probability that they succeed in coordinating on the same strategy in the underlying pure coordination game. Our first assumption, derived from more primitive assumptions on the monitoring technology in Section 6.3, registers the minimal requirements on the benefit function B consistent with its derivation from an underlying game of pure coordination.

Assumption 1

- (1.1) $B(\alpha, \beta) = B(\beta, \alpha)$.
- (1.2) $B(\alpha, \beta) \leq \min\{B(\alpha, \alpha), B(\beta, \beta)\}$.

The payoff or fitness derived from using α is the difference between the expected benefit and the *cost* $C(\alpha)$ of choosing α . The cost is paid regardless of whether any strategy in the pure coordination game actually turns out to have one of the monitored properties, or whether a realized property is actually used for anything after being observed. In thinking about $C(\alpha)$, we recognize that it may cost little to monitor a useful property like *biggest* and much to monitor a property like *anagram of a Brazilian poet's name*. We also recognize that increasing expenditure on monitoring a property like *largest* may reduce the cost of monitoring a property like *smallest*.

Our second assumption imposes some regularity conditions on the cost function C , and the function A defined by

$$A(\alpha) = B(\alpha, \alpha),$$

which is the benefit in the efficient case when both players use the same monitoring strategy. (See Section 6.3 again.)

Assumption 2

- (2.1) A is increasing and continuously differentiable.
- (2.2) C is increasing and continuously differentiable.

We let $A_u(\alpha)$ and $C_u(\alpha)$ denote the directional derivatives of the functions A and C in the direction of a unit vector u .

Assumption 3

- (3.1) C is strictly convex.
- (3.2) $A - C$ is strictly concave.

This assumption removes some obvious obstacles to the selection of an efficient equilibrium in the monitoring game. The concavity of $A - C$ ensures the existence of a unique efficient monitoring strategy. Section 6.4 illustrates how the failure of this assumption can trap evolution at an inefficient local optimum. Adding the convexity of C ensures that the “risk dominant” equilibrium is unique, simplifying some of the results below.

The benefit function $B(\alpha, \beta)$ need not be a concave function of β . However, it will sometimes be useful to make the following assumption.

Assumption 4 *Within the space of monitoring strategies, $B(\alpha, \beta)$ is a concave function of β on any line through $\beta = \alpha$.*

Although A and C are assumed to be differentiable, the next paragraph shows why the benefit function B can be differentiable at (α, α) only in pathological cases. However, left and right directional derivatives of B with respect to one of its variables commonly exist. Let $B_u^+(\alpha, \alpha)$ denote the “right” directional derivative, i.e. the derivative taken when moving away from α in the direction of u . The symmetry of B ensures that this derivative is the same, whether taken with respect to the first or second variable. Similarly, $B_u^-(\alpha, \alpha)$ is the “left” directional derivative, i.e. the derivative taken when moving in the direction of u towards α . If these directional derivatives exist and are continuous at (α, α) , we say B is *semi-differentiable*.

Observe that $B_u^-(\alpha, \alpha) \geq 0 \geq B_u^+(\alpha, \alpha)$ by Assumption 1.2. Section 6.5 shows that Assumption 4 and semi-differentiability imply

$$A_u(\alpha) = B_u^+(\alpha, \alpha) + B_u^-(\alpha, \alpha). \quad (1)$$

As a result, $B_u^-(\alpha, \alpha) = B_u^+(\alpha, \alpha)$ implies $A_u(\alpha) = 0$. The function $B(\alpha, \beta)$ therefore cannot normally be a differentiable function where $\beta = \alpha$.

The final assumption is needed only when considering stochastic potential. It says that the benefit sacrificed when deviating from α to β is greater when the other player is coordinating by also playing α than when he or she is miscoordinating by playing some other monitoring strategy γ .

Assumption 5 $B(\alpha, \alpha) - B(\beta, \alpha) \geq B(\alpha, \gamma) - B(\beta, \gamma)$.

The Square Monitoring Game. We associate a monitoring game with the Square Game of Section 2 by identifying a player’s monitoring intensity with the probability that he or she observes an odd-man-out when one exists. We assume that the players’ observations are *perfectly correlated*, by which we mean that a player who monitors a property with intensity α observes

whatever is observed by a player who monitors the same property with intensity $\beta \leq \alpha$, and perhaps more. (See Section 6.3.) We then have that $B(\alpha, \beta) = \frac{7}{8}\beta + (1 - \frac{7}{8}\beta)\frac{1}{3}$, when $\beta \leq \alpha$. Thus $A(\alpha) = \frac{1}{2}\alpha + \frac{1}{3}$, and

$$B(\alpha, \beta) = \begin{cases} A(\beta) & \text{if } \beta \leq \alpha \\ A(\alpha) & \text{if } \alpha \leq \beta \end{cases} . \quad (2)$$

Let the cost function be $C(\alpha) = \frac{1}{2}\alpha^2$. Then Assumptions 1–5 all hold. Also, B is semi-differentiable, with $B_u^-(\alpha, \alpha) = A'(\alpha) = \frac{1}{2}$, and $B_u^+(\alpha, \alpha) = 0$ when u points to the right.

The Red Square Monitoring Game. We associate a monitoring game with the Red Square Game of Section 2 by supposing that the monitoring strategy $\alpha = (\alpha_1, \alpha_2)$ results in a player identifying an odd-man-out in shape (when one exists) with probability α_1 , and an odd-man-out in color (when one exists) with probability α_2 . Again, players' observations are perfectly correlated. Assuming that both players use the Shape-then-Color convention, we have

$$B(\alpha, \beta) = \frac{1}{2}m_1 + \frac{1}{2}m_2(1 - \frac{3}{4}M_1) + \frac{1}{3},$$

where $m_i = \min\{\alpha_i, \beta_i\}$ and $M_i = \max\{\alpha_i, \beta_i\}$. Let $C(\alpha) = \frac{1}{4}(\alpha_1^2 + \alpha_2^2)$. Then Assumptions 1–5 hold. Also, B is semi-differentiable, with $B_u^+(\alpha, \alpha) = 0$ when $u = (1, 0)$.

4 Evolutionarily Stable Monitoring

4.1 Efficient equilibria.

The monitoring game is assumed to be unframed, and so only its symmetric equilibria are relevant.

A pure strategy α corresponds to a symmetric equilibrium when

$$A(\alpha) - C(\alpha) \geq B(\alpha, \beta) - C(\beta)$$

for all β . Assumption 1 ensures that this inequality always holds when $C(\beta) \geq C(\alpha)$. When $C(\beta) < C(\alpha)$, it is implied by the inequality $A(\beta) - A(\alpha) \leq C(\beta) - C(\alpha) \leq 0$. By Assumptions 2–3, the latter is equivalent to

$$0 \leq \frac{C_u(\alpha)}{A_u(\alpha)} \leq 1, \quad (3)$$

for all directional derivatives A_u and C_u for which u points into the space of monitoring strategies, and for which $C_u(\alpha) < 0$. As a result, there will typically be a multitude of equilibria. This is to be expected. Framing information is valuable only to the extent that it facilitates coordination. There are then many monitoring intensities that will be optimal if but only if they are chosen by one's fellow players. For example, all monitoring strategies α with $\alpha \leq \frac{1}{2}$ correspond to symmetric equilibria in the Square Monitoring Game.

Section 6.6 shows that efficient monitoring intensities must be pure. An efficient strategy α^* therefore maximizes $A(\alpha) - C(\alpha)$, and so satisfies $A_u(\alpha^*) = C_u(\alpha^*)$ for all u pointing into the space of monitoring strategies. Assumption (3) ensures that there is a unique efficient strategy.

The efficient monitoring strategy in the Square Monitoring Game is $\alpha^* = \frac{1}{2}$, which is the largest equilibrium strategy. More generally, (3) implies:

Lemma 1 *The efficient strategy α^* corresponds to an equilibrium of the monitoring game.*

Notice that α^* only just satisfies the second inequality in (3), which is our first indication that efficiency is evolutionarily precarious in monitoring games.

To avoid special cases, we assume that the efficient strategy α^* is interior, so that $A_u(\alpha^*) = C_u(\alpha^*)$ for all directions u . Every property is monitored with positive intensity at an interior monitoring strategy, and so will sometimes be used in identifying a focal point. When properties are monitored with low intensity, this is unlikely to be a conscious activity. As Schelling observed, one often does not anticipate what properties might make a strategy focal in unusual situations until confronted with the need to make a coordinating decision. For example, people commonly don't ask themselves whether going to Grand Central Station is one of their options until forced to consider possible meeting points in New York City.

4.2 Invasion barriers

Vickers and Canning [19] point out that it is empty to say that a strategy α is evolutionarily stable in an infinite game, unless it admits a positive *global invasion barrier* $\epsilon^*(\alpha)$. (See also Bomze and Pötscher [4].) This means that a mutant must take control of a fraction $\epsilon^*(\alpha)$ or more of a population using strategy α before it can invade under a monotonic evolutionary dynamic—no

matter what alternative strategy β the mutant adopts.⁷

A strategy's global invasion barrier provides a guide to how long the strategy is likely to survive before succumbing to a mutant invasion. The larger the barrier, the less likely are mutants to arrive in sufficient strength to overwhelm the existing strategy, and the longer can we expect the strategy to persist.

Definition 1 Let $\Pi(\alpha, (1-\epsilon)\alpha + \epsilon\beta)$ be the expected payoff to a player using monitoring strategy α in a population of which a fraction $1 - \epsilon$ use α , and a fraction ϵ use β . Then the global invasion barrier $\epsilon^*(\alpha)$ of α is

$$\epsilon^*(\alpha) = \sup\{\epsilon : \beta \neq \alpha \Rightarrow \Pi(\alpha, (1-\epsilon)\alpha + \epsilon\beta) > \Pi(\beta, (1-\epsilon)\alpha + \epsilon\beta)\}.$$

To simplify the notation we adopt the convention that attributing a negative value to $\epsilon^*(\alpha)$ is equivalent to asserting that $\epsilon^*(\alpha) = 0$. Attributing a value to $\epsilon^*(\alpha)$ that exceeds one is equivalent to asserting that $\epsilon^*(\alpha) = 1$.

Proposition 1 The global invasion barrier of α is

$$\epsilon^*(\alpha) = \inf_{\beta \neq \alpha} \left\{ \frac{C(\beta) - C(\alpha) - B(\alpha, \beta) + A(\alpha)}{A(\alpha) + A(\beta) - 2B(\alpha, \beta)} \right\}. \quad (4)$$

If B is semi-differentiable and satisfies Assumption 4, then

$$\epsilon^*(\alpha) = \inf_u \left\{ \frac{RB_u^- + (R-1)B_u^+}{B_u^- - B_u^+}, \frac{RB_u^+ + (R-1)B_u^-}{B_u^+ - B_u^-} \right\}, \quad (5)$$

where the infimum extends over all unit vectors u that point from α into the space of monitoring strategies, and where $R = C_u/A_u$, B_u^- and B_u^+ are all evaluated at (α, α) .

Proof. To find the fraction of a population using α that must be taken over by a mutant using β before the mutant can invade, we examine the reduced game of Figure 4. If

$$M(\alpha, \beta) = \frac{C(\beta) - C(\alpha) - B(\alpha, \beta) + A(\alpha)}{A(\alpha) + A(\beta) - 2B(\alpha, \beta)} \quad (6)$$

lies in $[0, 1]$, then the reduced game has a mixed equilibrium in which β is played with probability $M(\alpha, \beta)$. A figure like Figure 3 can then be drawn

⁷The standard definition of an ESS allows $\epsilon^*(\alpha)$ to be a function of β . This freedom is innocuous in finite games, but the infimum of all such ϵ may be zero in an infinite game.

	α	β
α	$A(\alpha) - C(\alpha)$	$B(\beta, \alpha) - C(\beta)$
β	$B(\alpha, \beta) - C(\alpha)$	$A(\beta) - C(\beta)$

Figure 4: The reduced monitoring game with strategies α and β .

in which the mixed equilibrium $M(\alpha, \beta)$ separates the regions where α and β are more fit. Hence, strategy β can invade if it initially seizes more than a fraction $M(\alpha, \beta)$ of the population. If $M(\alpha, \beta) > 1$ then α dominates β and the latter can never invade. If $M(\alpha, \beta) < 0$, then β is a superior reply to α , and the invasion barrier is zero, giving (4).

To obtain (5), take $\beta = \alpha + \rho u$ in (4). Assumption 4 ensures that $M(\alpha, \beta)$ decreases as β approaches α from any direction. Hence, it suffices to examine the limiting values of $M(\alpha, \beta)$ as $\rho \rightarrow 0+$ and $\rho \rightarrow 0-$, which are, respectively,

$$\frac{C_u(\alpha) - B_u^+(\alpha)}{A_u(\alpha) - 2B_u^+(\alpha)} \quad \text{and} \quad \frac{C_u(\alpha) - B_u^-(\alpha)}{A_u(\alpha) - 2B_u^-(\alpha)}. \quad (7)$$

If B is semi-differentiable, we can use (1) to rewrite the expressions in (7) to obtain those in (5). ||

The derivation of (5) exploits the fact that the mutant strategies that can most easily invade a population playing α lie arbitrarily close to α . This is the first of several indications that the mutants who are most likely to destabilize an equilibrium lie nearby in the space of monitoring strategies. This is consistent with the biological view of evolution. As Jacob [9] puts it, Nature tinkers by creating small variations on what she has already created rather than throwing “hopeful monsters” into the ring.

Corollary 1 *The monitoring strategy $\hat{\alpha}$ with the largest global invasion barrier is the unique α that maximizes*

$$A(\alpha) - 2C(\alpha),$$

It is always true that $\epsilon^(\hat{\alpha}) \geq \frac{1}{2}$. If B is semi-differentiable and satisfies Assumption 4, then $\epsilon^*(\hat{\alpha}) = \frac{1}{2}$.*

Proof. The expression (6) for $M(\alpha, \beta)$ is greater or less than $\frac{1}{2}$ depending on whether $B(\alpha) - 2C(\alpha)$ is greater or less than $B(\beta) - 2C(\beta)$. But $\hat{\alpha}$ is the unique monitoring strategy that maximizes $B(\alpha) - 2C(\alpha)$ (Assumption 3). Hence $\epsilon^*(\hat{\alpha}) \geq \frac{1}{2}$, and $\epsilon^*(\beta) < \frac{1}{2}$ for any $\beta \neq \alpha$. If B is semi-differentiable and Assumption 4 holds, then $B_u^+ + B_u^- = A_u = 2C_u$ at $\hat{\alpha}$ (by (1)). It follows that the two ratios in (5) both equal $\frac{1}{2}$, and hence so does $\epsilon^*(\hat{\alpha})$. \parallel

Corollary 2 *Suppose that B is semi-differentiable and satisfies the concavity condition of Assumption 4. Then the global invasion barrier of the efficient monitoring strategy α^* is zero if $B_u^+(\alpha^*, \alpha^*) = 0$ or $B_u^-(\alpha^*, \alpha^*) = 0$ for some direction u .*

Proof. It follows from (5) of Proposition 1 that

$$\epsilon^*(\alpha^*) \geq \min\{R, 1 - R\},$$

with equality when $B_u^+ = 0$ or $B_u^- = 0$ for some u . But $R = C_u(\alpha^*)/A_u(\alpha^*) = 1$, giving the result. \parallel

To interpret this conclusion, suppose that $A_u(\alpha^*, \alpha^*) > 0$ for some unit vector u . By (1), $B_u^+(\alpha^*, \alpha^*) + B_u^-(\alpha^*, \alpha^*) = A_u(\alpha^*)$, and so $B_u^+(\alpha^*, \alpha^*) = 0$ implies that $B_u^-(\alpha^*, \alpha^*) = A_u(\alpha^*, \alpha^*)$. A marginal reduction in the mutant's monitoring intensities below α^* (in the direction $-u$) thus has virtually no effect on the mutant's payoff against itself (an effect of approximately $A_u(\alpha^*) - C_u(\alpha^*) = 0$) or against the efficient strategy (an effect of approximately $B_u^-(\alpha^*, \alpha^*) - C_u(\alpha^*) = A_u(\alpha^*) - C_u(\alpha^*) = 0$). However, whenever it plays the mutant, this reduction imposes a loss of $B_u^-(\alpha^*, \alpha^*)$ that is accompanied by no corresponding cost savings for the existing strategy. The mutant thus earns a higher expected payoff and can invade. Again we see the importance of local mutants, and in this case the importance of mutants who monitor less intensely.

The Square Monitoring Game again. Inserting (2) in (6) and using the fact that A is concave and C convex, we have that

$$M(\alpha, \beta) = \begin{cases} 1 - \frac{C(\alpha) - C(\beta)}{A(\alpha) - A(\beta)} \geq 1 - \frac{C'(\alpha)}{A'(\alpha)} & \text{if } \beta < \alpha \\ \frac{C(\beta) - C(\alpha)}{A(\beta) - A(\alpha)} \geq \frac{C'(\alpha)}{A'(\alpha)} & \text{if } \beta > \alpha \end{cases}. \quad (8)$$

It follows that $M(\alpha, \beta)$ is smallest in the limit as $\beta \rightarrow \alpha$, and

$$\epsilon^*(\alpha) = \min \left\{ 1 - \frac{C'(\alpha)}{A'(\alpha)}, \frac{C'(\alpha)}{A'(\alpha)} \right\}. \quad (9)$$

The efficient strategy $\alpha^* = \frac{1}{2}$ has a global invasion barrier of zero. The largest global invasion barrier of $\frac{1}{2}$ occurs at $\hat{\alpha} = \frac{1}{4}$ (where $A'(\hat{\alpha}) = 2C'(\hat{\alpha})$). The efficient strategy α^* is therefore quite vulnerable to mutant invasion, whereas $\hat{\alpha}$ is relatively immune.

The Red Square Monitoring Game again. The efficient strategy $\alpha^* = (\frac{4}{7}, \frac{4}{7})$ has a global invasion barrier of zero. The largest global invasion barrier of $\frac{1}{2}$ occurs at $\hat{\alpha} = (\frac{4}{11}, \frac{4}{11})$.

Polymorphic invasions. We have focused throughout on invasions by waves of mutants who all play the same strategy. This is especially consistent with the biological literature, in which invasions are thought to originate with single mutations that become established in isolated subpopulations to which the parent population is subsequently exposed. Our results also hold with the polymorphic invasions, except that Corollary 1 then requires Assumption 5.

4.3 Risk Dominance

The condition that α risk dominates β in the reduced game of Figure 4 is that $M(\alpha, \beta) > \frac{1}{2}$ in (6). This holds if and only if

$$A(\alpha) - 2C(\alpha) > A(\beta) - 2C(\beta). \quad (10)$$

We can therefore characterize the monitoring strategy $\hat{\alpha}$ with the maximum global invasion barrier as the strategy that risk dominates all other strategies in pairwise comparisons. In particular:

Corollary 3 *The efficient monitoring strategy α^* is pairwise risk dominated.*

The efficient strategy α^* is not only risk-dominated in a pairwise comparison with $\hat{\alpha}$, but also by mutants lying arbitrarily close to, but monitoring less intensely than, α^* . The efficient strategy satisfies $A(\alpha^*) - A(\beta) \geq C(\alpha^*) - C(\beta)$, but β risk-dominates α^* if $A(\alpha^*) - A(\beta) < 2(C(\alpha^*) - C(\beta))$. For mutants lying sufficiently close to α^* , the first expression is nearly an equality, and so the second expression also holds when $C(\beta) < C(\alpha^*)$. We therefore gain some insight into how the efficient strategy may be destabilized by an invasion of less costly nearby strategies. At an efficient outcome the lost benefit in monitoring slightly less intensely is just balanced by the

decreased cost. Mutants who monitor a little less intensely therefore secure only slightly less than the efficient payoff when playing against copies of themselves, but they also achieve a saving in cost when playing against the efficient strategy, an advantage that allows lower-cost mutants to risk-dominate the efficient strategy in a pairwise comparison.

4.4 Stochastic Stability

We now examine a model of best-response dynamics with inertia in the spirit pioneered by Kandori, Mailath and Rob [10] and Young [20]. If M is the number of properties in a monitoring game, consider a grid of monitoring strategies in \mathbb{R}_+^M whose mesh is less than some small $\Delta > 0$. Because $B(\alpha) \leq 1$ and $C(\alpha)$ is increasing and convex, we can restrict attention to a finite subset of this grid, referred to as the set of *admissible* strategies, that excludes monitoring intensities which ensure negative payoffs. We assume without loss of generality that all strategies α of particular interest are points in the grid.

A large but finite population of η players play a coordination game in a round-robin tournament each period. At the end of each period, each player independently receives a signal with probability $\lambda > 0$. Players who receive the signal switch their strategy to a best response to the strategies played in the previous period. Other players continue to use their previous strategies. Each player is then independently selected to be a “mutant” with probability $\mu > 0$. Mutants independently switch to a new strategy that is equally likely to be any of those available in the grid.

These behavioral rules define the transition probabilities for a stochastic process. A state of this process is a specification of the number of players using each possible strategy in a monitoring game. The state α is to be understood as the state in which everyone plays α . We are interested in what happens in the ultralong run, when the mutation rates become vanishingly small. More precisely, we seek states that are *stochastically stable*, which means that they are assigned positive probability when we take the limit of the invariant distribution of our stochastic process as $\mu \rightarrow 0$ (Young [20], Kandori, Mailath and Rob [10], Samuelson [14]).

Proposition 2 *Let Assumption 5 hold. Then for sufficiently large η , the unique stochastically stable state is $\hat{\alpha}$, the monitoring strategy α that maximizes $A(\alpha) - 2C(\alpha)$.*

Proof. Strategy $\hat{\alpha}$ risk dominates every other pure strategy in a pairwise comparison and hence, for any pure strategy β in the finite set of admissible strategies, we have

$$\frac{1}{2}\pi(\hat{\alpha}, \hat{\alpha}) + \frac{1}{2}\pi(\hat{\alpha}, \beta) > \frac{1}{2}\pi(\beta, \hat{\alpha}) + \frac{1}{2}\pi(\beta, \beta). \quad (11)$$

Now suppose that half of the population plays $\hat{\alpha}$ and that proportions k_1, \dots, k_m of the population play strategies $\gamma_1, \dots, \gamma_m$, respectively. We then note that, for any admissible β :

$$\begin{aligned} \frac{1}{2}\pi(\hat{\alpha}, \hat{\alpha}) + \sum_{i=1}^m k_i \pi(\hat{\alpha}, \gamma_i) &= \frac{1}{2}\pi(\hat{\alpha}, \hat{\alpha}) + \sum_{i=1}^m k_i \pi(\hat{\alpha}, \beta) + \sum_{i=1}^m k_i [B(\hat{\alpha}, \gamma_i) - B(\hat{\alpha}, \beta)] \\ &> \frac{1}{2}\pi(\beta, \hat{\alpha}) + \sum_{i=1}^m k_i \pi(\beta, \beta) + \sum_{i=1}^m k_i [B(\beta, \gamma_i) - B(\beta, \beta)] \\ &= \frac{1}{2}\pi(\beta, \hat{\alpha}) + \sum_{i=1}^m k_i \pi(\beta, \gamma_i), \end{aligned}$$

where the inequality follows from (11) and (5) (which, by a suitable permutation of α , β , and γ , gives $B(\hat{\alpha}, \gamma_i) - B(\hat{\alpha}, \beta) > B(\beta, \hat{\alpha}) - B(\beta, \beta)$). We have then shown that $\hat{\alpha}$ is a strict best response to any population in which at least half of the agents play $\hat{\alpha}$, i.e., it is $\frac{1}{2}$ -dominant. Ellison shows that $\hat{\alpha}$ must then be the unique stochastically stable state ([7, Corollary 1]). \parallel

When there is a unique stochastically stable state, it is the root of the “minimum-resistance” tree, where the nodes of such a tree are the states of the system, and each edge carries a resistance which measures the minimum number of mutations required to get the system from the origin to the terminus of the edge, using the best-response dynamics to hitch a free ride where possible. We illustrate this construction in the case of the Square Monitoring Game.

The Square Monitoring Game yet again. In considering the resistance in passing from state α to state β , the location of the mixed equilibrium of the reduced monitoring game containing only strategies α and β is relevant. If this requires that β be played with probability $M(\alpha, \beta)$, then we need $[\eta M(\alpha, \beta)] + 1$ simultaneous mutations to bounce the system from α into the basin of attraction of β . When the population size η is large, it follows that the global invasion barrier of α identifies the minimum resistance in passing from α to some other equilibrium β , which we say is the “nearest” equilibrium to α .




Figure 5: Minimizing stochastic potential in the Square Monitoring Game.

It follows from (8) that the global invasion barrier of α is $C'(\alpha)/A'(\alpha)$ when $A'(\alpha) - 2C'(\alpha) \geq 0$, and the “nearest” equilibrium lies immediately to the right of α . If $A'(\alpha) - 2C'(\alpha) \leq 0$, then the global invasion barrier is $1 - C'(\alpha)/A'(\alpha)$, and the “nearest” equilibrium lies immediately to the left.

We construct the tree shown in Figure 5 by joining each equilibrium other than $\hat{\alpha}$ to its “nearest” neighbor. Pending attention to two complications, this tree minimizes total resistance. First, we have included as nodes in our tree only states that correspond to monomorphic equilibria. However, the best response dynamics allow us to move from any nonequilibrium state of this game to an equilibrium without incurring any mutations, and hence without increasing the resistance of the tree, so that the latter states can be neglected. Secondly, we have assumed that when making a transition from state α to β , the minimum number of mutations can be achieved by having all mutants play strategy β . Assumption 5 ensures that this is the case. Hence, $\hat{\alpha}$ is the unique stochastically stable equilibrium of the Square Monitoring Game.

The Red-Square Monitoring Game yet again. A similar argument allows us to draw a schematic of the minimal tree of Figure 6, whose root lies at the inefficient equilibrium $\hat{\alpha} = (\frac{4}{11}, \frac{4}{11})$.

5 Discussion

In real-life games, focal or salience considerations are often used to choose between multiple equilibria. Schelling’s [16] work shows that the process by which we convert cues from the environment into a choice of a coordination strategy is largely unconscious. It is only when confronted with his formal examples—in which people are clear about what is focal without being able




Figure 6: The minimal tree for the Red-Square Monitoring Game. The functions are given by $\alpha_2/(1 - \frac{3}{4}\alpha_1) = (1 - \alpha_1)/(1 + \frac{3}{4}\alpha_2)$ and $\alpha_2/(1 - \frac{3}{4}\alpha_1) = 1 - (1 - \alpha_1)/(1 + \frac{3}{4}\alpha_2)$.

to say why—that one realizes that we must be continuously monitoring a whole range of possible coordinating cues from the environment. We call the collection of such contextual cues a framing of a real-life game.

Game theorists have traditionally abstracted away framing effects in order to construct a world in which the question of what constitutes rational behavior can be studied without distraction. This paper moves in the other direction by incorporating the choice of how much attention players choose to pay to the frame of a game as a move in an evolutionary game within which the pure coordination game of traditional game theory is embedded. The result is a game in which evolutionary forces will tend to select inefficiently low monitoring intensities.

The intuition is simple. Suppose that a population using inefficiently low monitoring intensities is invaded by a mutant who monitors more intensely. The higher monitoring costs may be more than compensated for by the increased frequency of coordination when the mutant is matched with a copy of itself, but the mutant continues to bear the higher monitoring costs when playing against the incumbent strategy without any compensating increase in benefit. As a result, increased monitoring that would be efficient if performed by the whole population need not enjoy evolutionary success. Conversely, efficient outcomes are vulnerable to invasion by a succession of nearby mutants who monitor less intensely.

Under relatively mild conditions, we argue that one is likely to find a population at an equilibrium where marginal benefit equals *twice* marginal cost—rather than at the efficient equilibrium where marginal benefit equals *once* marginal cost—provided that the available time span is sufficiently long. The *two* in this result appears for essentially the same reason that one expects a stochastic evolutionary process to take a population to the risk-dominant equilibrium rather than the payoff-dominant equilibrium in a 2×2 game of impure coordination.

6 Appendix

6.1 One Independent Property

Let $V(N, p)$ be the expected payoff in a game with N strategies and a single property that is independently assigned to each strategy with probability p .

Lemma 2 *Let $p^*(N)$ maximize $V(N, p)$. Then*

$$\lim_{N \rightarrow \infty} p^*(N) \approx \frac{1.30}{N}$$

$$\lim_{N \rightarrow \infty} V(N, p^*(N)) \approx .50.$$

Proof. We have $V(N, p) =$

$$\left(p^N + (1-p)^N\right) \frac{1}{N} + \sum_{k=1}^{N-1} \binom{N}{k} \left(p^k (1-p)^{N-k}\right) \max\left\{\frac{1}{k}, \frac{1}{N-k}\right\}. \quad (12)$$

The first term captures the event that the property either assigns a 0 to each strategy or assigns a 1 to each strategy, providing no help with coordination and dooming the players to a payoff of $1/N$. The summation captures those cases in which at least one strategy is assigned a 1 and at least one strategy is assigned a 0. The maximum appears in the final term because players facing k 1s can either coordinate on the 1s to achieve a payoff of $1/k$ or coordinate on the complementary 0s to achieve a payoff of $1/(N-k)$, and will do whichever produces the higher payoff.

For each integer $N \geq 1$, let p_N maximize $V(N, p)$. The probability of precisely k ones being drawn in N independent trials is

$$\binom{N}{k} p_N^k (1-p_N)^{N-k} = p_N N k \binom{N-1}{k-1} p_N^{k-1} (1-p_N)^{(N-1)-(k-1)}. \quad (13)$$

Now suppose that $Np_N \rightarrow c$ as $N \rightarrow \infty$ through some subsequence and for some constant c . If $c = \infty$, then the right side of (13) can remain finite (as it must, since the left side is a probability) only if the probability of $k - 1$ 1s approaches zero as N grows, for any $k \geq 1$. But then the expected payoff achieved by coordinating on strategies assigned a 1 approaches zero. If $c = 0$, then the right side of (13) approaches zero, requiring the probability of k 1s (the left side) to approach zero for any $k \geq 1$, again giving a zero payoff from coordinating on strategies assigned a 1. Applying a symmetric argument to the expected number of zeros $(1 - p_N)N$, we find that optimality requires that either $P_N N$ or $(1 - p_N)N$ approach a positive, finite constant. Without loss of generality, assume $P_N N \rightarrow c$ with $0 < c < \infty$.

Now, observe that, for a fixed value of $p_N N = c$,

$$\begin{aligned} \binom{N}{k} p_N^k (1 - p_N)^{N-k} &= \binom{N}{k} \left(\frac{c}{N-c}\right)^k \left(1 - \frac{c}{N}\right)^{N-k} \\ &= \frac{N}{N-c} \frac{N-1}{N-c} \times \cdots \times \frac{N-k+1}{N-c} \frac{c^k}{k!} \left(1 - \frac{c}{N}\right)^N \\ &\rightarrow \frac{c^k}{k!} e^{-c}. \end{aligned} \tag{14}$$

Hence, in the limiting case of large N , the likely number of 1s is governed by a Poisson distribution with mean c . Since an observation of no 1s gives an expected payoff of $1/N$, and an observation of $k > 0$ 1s gives an expected payoff of $1/k$, the limiting payoff (12), as $N \rightarrow \infty$, is given by

$$\sum_{k=1}^{\infty} \frac{1}{k} \frac{c^k}{k!} e^{-c}. \tag{15}$$

We now note that the expression in (15) has a unique local maximum. In particular, $c = 0$ is a local minimum, while a zero first derivative of (15) implies that the second derivative is negative. Next, we can find sufficiently large \bar{c} and \bar{N} such that for all $c > \bar{c}$ and $N > \bar{N}$,

$$V(N, c/N) < V(N, 1/N).$$

This follows from the fact that for $p_N N = 1$, the probability of exactly one value of 1 approaches $1/e$, and hence the payoff $V(N, 1/N)$ has a lower bound arbitrarily close to $1/e$ for large N . In contrast, for sufficiently large values of $p_N N = c$, the probability of fewer than $c/2$ or more than $2c$ ones is also arbitrarily small when N is large, ensuring that for large c , $V(N, c/N)$

has a limiting upper bound arbitrarily close to $2/c$, which suffices for the claim. Hence, we can restrict attention to values of c in an interval $[0, \bar{c}]$.

We now note that (12) converges to (15) uniformly on the compact interval $[0, \bar{c}]$. It suffices for this conclusion that the convergence in (14) is uniform for any k . To establish the latter, we note that $(1 - \frac{c}{N})^N$ is a decreasing sequence of functions converging to the continuous limit e^{-c} on the compact set $[0, \bar{c}]$, and hence converges uniformly (Rudin [13, Theorem 7.13]). We next note that $\frac{N}{N-c} \frac{N-1}{N-c} \times \dots \times \frac{N-k+1}{N-c}$ is minimized at $c = 0$ and maximized at $c = \bar{c}$, and hence furthest from its limit (as N grows) of unity for one of $c \in \{0, \bar{c}\}$. The convergence for these two values of c thus implies uniform convergence. Together, these imply the uniform convergence of (12) to (15).

Any maximizer of (12) must then converge to the maximizer of (15). As N grows, the optimal value of $p_N N$ thus approaches the latter. Solving the maximization problem given by (15) numerically, we have an optimal value of $c \approx 1.30$ and hence

$$p^*(N) \rightarrow \frac{1.30}{N}. \quad (16)$$

Similarly, we can solve numerically for the limiting optimal payoff (equal to approximately .50). ||

6.2 Many Independent Properties

Let there be M properties P_1, \dots, P_M , each of which is independently assigned to each strategy with a probability p_i that is constant across strategies for a given property but may vary across properties. Let c^* be the optimal value of pN when there is a single such strategy (≈ 1.30), calculated in the previous lemma. Let $V(N, M)$ be the expected payoff from a game of size N with M optimally-chosen such properties, and let $V(M) = \lim_{N \rightarrow \infty} V(N, M)$. Then:

Lemma 3

$$V(M) \geq 1 - \left(1 - \frac{1}{e}\right)^{2^{M-1}} \quad (17)$$

$$V(M+1) - V(M) \geq (1 - V(M)) \frac{1}{e}. \quad (18)$$

Proof. First, let $p_i = \frac{1}{2}$ for $i = 1, \dots, n-1$ and $p_n = 2^{M-1}/N$. Then as N gets large, each of the first $M-1$ properties produces a proportion of 1s arbitrarily close to $N/2$. Hence, we can proceed as if the strategies

are divided into 2^{M-1} groups of size $N/2^{M-1}$ each. With a probability that approaches $1/e$, the M th property produces precisely one 1 on such a group, allowing a payoff of one. Underestimating the expected payoff by assuming the payoff to be zero whenever we do not achieve at least one group with exactly one 1, we obtain (17), which can be reformulated to give (18). \parallel

It is straightforward to show that this configuration of properties, with the exception that p_M is chosen optimally, is payoff-maximizing, so that $V(M)$ is maximized by letting

$$p_i = \frac{1}{2}, \quad i = 1, \dots, M-1, \quad p_M = \frac{c^* 2^{M-1}}{N}.$$

6.3 The Benefit Function

Let \mathcal{P} and \mathcal{P}' be subsets of $\{P_1, P_2, \dots, P_M\}$ and let $H(\mathcal{P}, \mathcal{P}')$ be the probability of successful coordination, and hence expected payoff, when two players observe the set of properties \mathcal{P} and \mathcal{P}' , respectively. Recall that observing the set \mathcal{P} ensures that a player observes any realizations of these properties, but does not ensure that there actually are any realizations.⁸ It follows immediately from the symmetry of the coordination game that

$$H(\mathcal{P}, \mathcal{P}') = H(\mathcal{P}', \mathcal{P}).$$

In addition, we have

$$H(\mathcal{P}, \mathcal{P}') \leq H(\mathcal{P}, \mathcal{P}).$$

In particular, consider player i facing an opponent j who observes set \mathcal{P} . Observing \mathcal{P} provides player i with complete information about j 's observation and hence equilibrium behavior. Observing \mathcal{P}' instead of \mathcal{P} can only provide less information, and hence the (weak) inequality.

For each subset \mathcal{P} of $\{P_1, P_2, \dots, P_M\}$, we let $G(\mathcal{P}, \alpha)$ be the probability that a player using the monitoring strategy α observes the realizations of all of the properties in \mathcal{P} and of no properties not in \mathcal{P} . Again, some of the observed properties in \mathcal{P} will (disappointingly) be held by none of the strategies, while others may hold for many strategies; at the same time that

⁸To keep the notation simple, we assume that the realization of a property for one strategy is recognized if and only if all realizations of that property are recognized. This assumption sacrifices no generality (given our assumption below concerning correlation across players). Assuming that players may perceive some realizations and not others is equivalent to assuming that all realizations are noticed (if any are) but adjusting the probability with which realizations occur.

some of the unobserved properties will be exhibited by many strategies and some by none.

We assume that $G(\mathcal{P}, \alpha)$ is differentiable and make two monotonicity assumptions. First, for each set of properties \mathcal{P} with $P_i \in \mathcal{P}$

$$\frac{dG(\mathcal{P}, \alpha)}{d\alpha_i} > 0. \quad (19)$$

Hence, monitoring property P_i more intensely increases the probability that a set of properties including P_i is observed. Next, for any nonempty set \mathcal{P} with $P_i \notin \mathcal{P}$,

$$\frac{dG(\mathcal{P} \cup \{P_i\}, \alpha)}{d\alpha_i} > \frac{dG(\mathcal{P}, \alpha)}{d\alpha_i} \geq - \left| \frac{dG(\mathcal{P} \cup \{P_i\}, \alpha)}{d\alpha_i} \right|. \quad (20)$$

Hence, monitoring property i more intensely may decrease the probability that a set excluding P_i is observed, but not at a rate that exceeds the effect of the increased intensity on sets including P_i . This latter stipulation ensures that increasing one's monitoring intensity cannot decrease the total probability that *some* property is observed.⁹

For example, in the case of M independent properties, it is natural (but not necessary) to assume that there are M functions $g_1(\alpha_1), \dots, g_m(\alpha_m)$ from \mathbb{R}_+ into $[0, 1]$ such that property P_i is observed with probability $g_i(\alpha_i)$, where the realizations of these random variables are independent across properties. We then have

$$G(\mathcal{P}, \alpha) = \prod_{P_i \in \mathcal{P}} g(\alpha_i) \prod_{P_i \notin \mathcal{P}} (1 - g(\alpha_i))$$

and hence

$$\begin{aligned} \frac{dG(\mathcal{P} \cup \{P_i\})}{d\alpha_i} &= \frac{dg(\alpha_i)}{d\alpha_i} \prod_{P_j \in \mathcal{P}} g(\alpha_j) \prod_{P_j \notin \{\mathcal{P} \cup \{P_i\}\}} (1 - g(\alpha_j)) \\ \frac{dG(\mathcal{P} \cup \{P_i\})}{d\alpha_i} &= - \frac{dg(\alpha_i)}{d\alpha_i} \prod_{P_j \in \mathcal{P}} g(\alpha_j) \prod_{P_j \notin \{\mathcal{P} \cup \{P_i\}\}} (1 - g(\alpha_j)), \end{aligned}$$

⁹The change in the probability that some property is observed is given by

$$\sum_{\mathcal{P} \neq \emptyset: P_i \notin \mathcal{P}} \frac{dG(\mathcal{P} \cup \{P_i\}, \alpha)}{d\alpha_i} + \sum_{\mathcal{P} \neq \emptyset: P_i \in \mathcal{P}} \frac{dG(\mathcal{P}, \alpha)}{d\alpha_i},$$

which (20) ensures is nonnegative.

ensuring that (20) holds, with equality in the second weak inequality. More generally, condition (20) also accommodates dependencies between properties.

We let $G(\mathcal{P}', \alpha' | \mathcal{P}, \alpha)$ be the probability that a player whose monitoring intensity is α' observes set \mathcal{P}' given that a player with intensities α observes \mathcal{P} . At one extreme, we could imagine the players' monitoring outcomes being independent, so that $G(\mathcal{P}', \alpha' | \mathcal{P}, \alpha) = G(\mathcal{P}', \alpha')$. In general, however, we expect such independence to fail. Instead, we think of Nature as determining which strategies have which properties and also determining how obvious these properties are. In some cases, there may be such an obvious central strategy that all players observe this property, even those who monitor "central" with a low intensity. In other cases, one of the strategies may be interpreted as central, but only after some reflection, so that only those who monitor "central" with a sufficiently high intensity observe the property. The important consideration is that we expect observations of "central" by players who monitor this property with high intensity to be correlated with observations on the part of those who monitor the property with lower intensity. We capture this by assuming that monitoring is *perfectly correlated across players*: if player i 's monitoring intensities make i more likely than j to observe a set of properties \mathcal{P} , then j observes the set \mathcal{P} only if i does. Hence, more vigilant players are always no less well informed than less vigilant players. Nature determines which strategies have which properties and also determines how obvious these properties are, with a set of properties then being observed by every player whose monitoring makes that set sufficiently likely.¹⁰

A player choosing monitoring intensities α and facing an opponent choosing intensities α' has an expected benefit given by the function $B(\alpha, \alpha')$ and an expected fitness given by

$$\begin{aligned} \pi(\alpha, \alpha') &= B(\alpha, \alpha') - C(\alpha) \\ &= \sum_{\mathcal{P}} \sum_{\mathcal{P}'} H(\mathcal{P}, \mathcal{P}') G(\mathcal{P}, \alpha | \mathcal{P}', \alpha') G(\mathcal{P}', \alpha') - C(\alpha). \end{aligned}$$

¹⁰It is straightforward to construct formally an underlying probability space with these properties. For example, suppose that there is only one property P_1 , so that the two relevant sets are $\{P_1\}$ and \emptyset . Let the monitoring intensities cause player 1 to observe $\{P_1\}$ with probability .4 and player 2 with probability .6. To achieve perfectly correlated monitoring, we need only draw a realization of a random variable z that is uniformly distributed on $[0, 1]$, with both players observing P_1 when $z \leq .4$, only player 2 when $p \in (.4, .6]$, and neither player when $p > .6$. Cases with more properties can be addressed with more elaborate but analogous constructions.

It follows immediately from the symmetry of the coordination game that

$$B(\alpha, \beta) = B(\beta, \alpha).$$

The assumption that monitoring is perfectly correlated across players ensures that

$$B(\alpha, \beta) \leq \min\{B(\alpha, \alpha), B(\beta, \beta)\},$$

since monitoring that is perfectly correlated across players ensures that when player i plays intensities α against an opponent playing α , that i 's payoff conditional on j observing \mathcal{P} is $H(\mathcal{P}, \mathcal{P})$. If instead i plays α' , then i 's payoff conditional on j observing \mathcal{P} can only be lower. This establishes Assumption 1.

The continuous differentiability of $A(\alpha)$ follows from the corresponding differentiability assumption on $G(\mathcal{P}', \alpha' | \mathcal{P}, \alpha)$.

Lemma 4 $A(\alpha)$ is increasing in α .

Proof. We have

$$A(\alpha) = \sum_{\mathcal{P}} H(\mathcal{P}) G(\mathcal{P}, \alpha) = \sum_{\mathcal{P}} \left(\frac{h(0, \mathcal{P})}{N} + \sum_{k=1}^N \frac{h(k, \mathcal{P})}{k} \right) G(\mathcal{P}, \alpha),$$

where $h(k, \mathcal{P})$ is the probability that the minimum number of strategies distinguished by a common joint property is k after observing \mathcal{P} .¹¹ Let $\Psi(i)$ be the set of sets of properties that do not include property i . Then we have:

$$\frac{dA(\alpha)}{d\alpha_i} = \sum_{\mathcal{P} \notin \Psi(i)} H(\mathcal{P}) \frac{dG(\mathcal{P}, \alpha)}{d\alpha_i} + \sum_{\mathcal{P} \in \Psi(i)} H(\mathcal{P} \cup \{i\}) \frac{dG(\mathcal{P} \cup \{i\}, \alpha)}{d\alpha_i}.$$

Because $H(\mathcal{P} \cup \{i\}) \geq H(\mathcal{P})$, it follows from (20) that $dA(\alpha)/d\alpha_i \geq 0$. \parallel

6.4 Nonconcave $A(\alpha) - C(\alpha)$

Let there be a single property P . Let $g(\alpha)$ be the probability that this property is observed by a player who monitors with intensity $\alpha \in \mathbb{R}^+$, where

$$g(\alpha) = 1 - \frac{1}{e^{\alpha^2}}. \quad (21)$$

¹¹Notice that $h(k, \mathcal{P}) = 0$ when $N/2 < k < N$, since the complementary set of strategies must then contain a smaller nonempty set characterized by a common joint property.




Figure 7: Benefit and cost functions for which $A(\alpha) - C(\alpha)$ is not concave.

Let $H(P, P) = .5$ and $C(\alpha) = .3\alpha$. Let N be sufficiently large that we can take $1/N \approx 0$. Figure 7 shows the function

$$A(\alpha) = \frac{1}{2} \left(1 - \frac{1}{e^{\alpha^2}} \right) \quad (22)$$

as well as the cost function. It is immaterial that the cost function is only weakly concave, but significant that $A(\alpha) - C(\alpha)$ is not concave. The fitness $\pi(\alpha) = A(\alpha) - C(\alpha)$ has two local maxima, at $\alpha = 0$ and $\alpha \approx 1.16$, separated by a local minimum at $\alpha \approx .34$. It is then easy to construct evolutionary processes that can become trapped at the inefficient equilibrium $\alpha = 0$. It would be common to say that this “evolutionary landscape is rugged,” in which case it is no surprise that evolution can lead to inefficient outcomes.

In the case of a single property P , the concavity of $A(\alpha) - C(\alpha)$ can be ensured by assuming that $C(\alpha)$ is convex and the function $G(\{P\}, \alpha)$ is concave (in contrast to (21)). In more complicated settings, the concavity of $A(\alpha) - C(\alpha)$ becomes a joint restriction on the monitoring technology.

6.5 Semi-Differentiability

Let B be semidifferentiable. Then notice that

$$\frac{B(\alpha + \rho u, \alpha + \rho u) - B(\alpha, \alpha)}{\rho u} = \frac{B(\alpha + \rho u, \alpha + \rho u) - B(\alpha + \rho u, \alpha)}{\rho u} + \frac{B(\alpha + \rho u, \alpha) - B(\alpha, \alpha)}{\rho u}. \quad (23)$$

As $\rho \downarrow 0$, the left side converges to $A_u(\alpha)$. The second term on the left side converges to $B_u^+(\alpha, \alpha)$. Rewrite the first term on the left side as

$$\frac{B(\beta, \beta - \rho u) - B(\beta, \beta)}{-\rho u}, \quad (24)$$

where $\beta = \alpha + \rho u$. Now consider a compact set $\mathcal{B}(\alpha)$ of values of β containing α in its interior. For a fixed β in this set, (24) converges to $B_u^-(\beta, \beta)$ as $\rho \downarrow 0$. Moreover, Assumption 4 ensures that (24) is an increasing sequence. Rudin [13, Theorem 7.13] (replacing f_n with $-f_n$ in the statement of the theorem) then implies that the convergence is uniform on $\mathcal{B}(\alpha)$. Using the uniformity of the convergence and the continuity of $B_u^-(\beta, \beta)$, for any $\epsilon > 0$, there must be a $\rho(\epsilon)$ sufficiently small that (24) is within $\epsilon/2$ of $B_u^-(\beta, \beta)$ for any $\beta \in \mathcal{B}(\alpha)$, and also such that $B_u^-(\alpha + \rho u, \alpha + \rho u)$ is within $\epsilon/2$ of $B_u^-(\alpha, \alpha)$. But then the first term on the left side of (23) must converge to $B_u^-(\alpha, \alpha)$, giving the result.

6.6 Efficient Monitoring

Lemma 5 *An efficient outcome can always be achieved with pure, symmetric monitoring intensities.*

Proof. We first recall that

$$B(\alpha, \alpha) \geq B(\beta, \alpha). \quad (25)$$

Hence, the highest payoff realization in any efficient mixed strategy must be of the form $B(\alpha, \alpha) - C(\alpha)$ for some α in the support of the mixed strategy. The pure strategy α must then either be efficient, or payoffs must be constant over all realized outcomes in the mixed strategy. The latter is precluded by Assumption 3. \parallel

References

- [1] Michael Bacharach. Variable universe games. In Ken Binmore, Alan Kirman, and Piero Tani, editors, *Frontiers of Game Theory*, pages 255–276. MIT Press, Cambridge, Massachusetts, 1993.

- [2] Michael Bacharach and Dale O. Stahl. Variable-frame level- n theory. *Games and Economic Behavior*, 32:220–246, 2000.
- [3] Ken Binmore, Larry Samuelson, and Richard Vaughan. Musical chairs: Modelling noisy evolution. *Games and Economic Behavior*, 11:1–35, 1995.
- [4] Immanuel M. Bomze and Benedikt M. Pötscher. *Game Theoretical Foundations of Evolutionary Stability*. Springer Verlag, Berlin, 1989.
- [5] André Casajus. Focal points in framed strategic forms. *Games and Economic Behavior*, 32:263–291, 2000.
- [6] Vincent P. Crawford and Hans Haller. Learning how to cooperate: Optimal play in repeated coordination games. *Econometrica*, 58:571, 1990.
- [7] Glenn Ellison. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–46, 2000.
- [8] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, Massachusetts, 1988.
- [9] Francois Jacob. Molecular tinkering in evolution. In D. S. Bendall, editor, *Evolution from Molecules to Men*, pages 131–144. Cambridge University Press, Cambridge, 1983.
- [10] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.
- [11] John Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.
- [12] Arthur J. Robson and Fernando Vega-Redondo. Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory*, 70:65–92, 1996.
- [13] Walter Rudin. *Principles of Mathematical Analysis*. McGraw Hill, New York, 1976.
- [14] Larry Samuelson. Stochastic stability in games with alternative best replies. *Journal of Economic Theory*, 64:35–65, 1994.

- [15] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391, 1992.
- [16] Thomas Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA., 1980 (first edition 1960).
- [17] Reinhard Selten. A note on evolutionarily stable strategies in asymmetric animal contests. *Journal of Theoretical Biology*, 84:93–101, 1980.
- [18] Robert Sugden. A theory of focal points. *Economic Journal*, 105:533–550, 1995.
- [19] G. T. Vickers and C. Cannings. On the definition of an evolutionarily stable strategy. *Journal of Theoretical Biology*, 129:349–353, 1987.
- [20] Peyton Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.
- [21] Peyton Young. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, 1998.