

THE ASSESSMENT: GAMES AND COORDINATION

DAVID P. MYATT

University of Oxford

HYUN SONG SHIN

London School of Economics

CHRIS WALLACE

*University of Oxford*¹

Coordination problems arise in a multitude of economic interactions. Recent advances in the field of game theory have shed new light on these problems and the ways in which they might be analysed. This issue of the Oxford Review of Economic Policy first examines some of the theoretical dimensions to this literature, as well as some empirical and experimental insights. It goes on to apply some of these ideas to a number of important policy areas, including macroeconomic policy coordination, public good provision, and problems of political coordination.

I. GAME THEORY IN THE SOCIAL SCIENCES

During the latter half of the twentieth century, game theory rose to become a central feature of modern microeconomic analysis.² Prior to the advent of game-theoretic methods, classical economics focused upon the analysis of price-driven market systems. In these systems, economic agents are

faced with a set of market prices containing all the relevant information for their decisions. Taking prices as given, each economic agent acts optimally. This idea of simultaneous and mutually compatible optimization is captured by the notion of a competitive equilibrium, and is expounded in a research programme which culminated in the contributions of Arrow and Debreu.³ A competitive equilibrium comprises a set of prices and transactions such that

¹ The authors would like to thank Chris Allsopp, Alison Gomm, Ken Mayhew, and the contributing authors to this special issue for their help in the preparation of this paper.

² Game theory is an integral part of many modern microeconomics textbooks, such as Kreps (1990) and Mas-Colell *et al.* (1995).

³ For the full exposition see, for instance, Arrow and Debreu (1954), Debreu (1959), and Arrow and Hahn (1971).

each agent's transaction is feasible and optimal given those prices.

Of course, there were notable exceptions to this rule—among them were the classic contributions of Cournot (1838), Bertrand (1883), and Edgeworth (1897), each of whom considered the oligopolistic interaction of firms. Nowadays it is common to view all of these contributions through the lens of *game theory*: the study of strategic interaction. Strategic interactions arise when an individual agent's decision has the potential to affect the decision another might wish to take. Notice that this does not generally require the presence of a price system. The ideas of competitive equilibrium fit well with the textbook (Robbins, 1932) view of economics as the study of 'scarce resources versus infinite wants'.⁴ Game theory allows a broader conception of economics. In his assessment of John Nash's contributions to economic theory, Myerson (1999, 2002) argued that:

To understand the importance of Nash's work . . . we should begin with the very definition of economics itself. A generation before Nash could have accepted a narrower definition of economics, as a specialized social science concerned with the production and allocation of material goods . . . today economists can define their field more broadly, as being about the analysis of incentives in all social institutions.

So what precisely was the contribution of Nash and the other early game theorists?⁵ First, Nash (1950*b*, 1951), von Neumann and Morgenstern (1967), and others provided a language within which we may formulate strategic decision-making problems. An interaction can be represented as a *normal-form* or *strategic-form* game. Each strategic-form representation has at least three elements: players, their available strategies (or actions), and the payoffs they receive given any combination of strategies

chosen by themselves and the other agents in the game. Second, Nash (1950*a*, 1951) provided a solution concept—now known as *Nash equilibrium*. A Nash equilibrium specifies a set of strategies (one for each player) such that each player's strategy is a best response to the strategy associated with each of the other players.

Notice that a Nash equilibrium shares some of the features of a competitive equilibrium. In particular, no player in a Nash equilibrium has an incentive to deviate from her specified strategy. Nor does any agent in a competitive equilibrium have an incentive to alter the quantities involved in her transactions. Furthermore, in a competitive equilibrium, each agent faces the same set of prices, while a Nash equilibrium might be seen as a commonly shared expectation of how the game will be played. Hence both concepts admit, in some sense, an 'optimizing behaviour with rational expectations' interpretation.

If we allow the representation of a given interaction as a game, would those involved in fact play a Nash equilibrium? This question must be asked even when a Nash equilibrium is unique. Many games of interest, however, possess multiple Nash equilibria. Settings of interest (some of which are discussed in section II and other papers in this issue) include technology choice, speculative currency attacks, bank runs, wage bargaining, and tactical voting. When games exhibit multiple equilibria an additional question arises—which one (if any) will be played? Notice that these questions are often asked of competitive equilibria. Interestingly, the answers are not unrelated.⁶

A possible justification for Nash equilibrium play might involve the imposition of some degree of 'rationality' on the part of the agents involved.⁷ For

⁴ More accurately, Robbins (1932) describes economics as 'the science which studies human behaviour as a relationship between scarce means which have alternative uses'.

⁵ Some of the classic contributions are collected together by Kuhn (1997).

⁶ The paper by Jeffery Amato, Stephen Morris, and Hyun Shin in this issue employs a competitive model coupled with an informational specification similar to that present in the games studied here.

⁷ This is by no means the only possible justification for the play of a Nash equilibrium. Alternative avenues of research have been pursued by various authors. Biologists introduce the concept of evolutionarily stable strategies (ESS), which is closely related to Nash equilibrium (see Maynard Smith, 1974). They study the way in which evolutionary forces might result in the play of an ESS (for a good survey of this material see Hofbauer and Sigmund, 1988). Similar evolutionary justifications can be provided for the play of Nash equilibria (see Weibull (1995) for a very complete study). Another (closely related) approach is the stochastic adjustment dynamics literature, expositions of which can be found in Young (1998), Samuelson (1997), and Vega-Redondo (1996). The fictitious play literature takes yet a different route toward the same goal. See Fudenberg and Levine (1998) for a description of these ideas, along with a general treatment of learning to play Nash equilibria in games.

example, players might be assumed to have consistent preferences over outcomes which guide their choice. As a result they might be assumed to choose the best alternative given the beliefs they hold. This, however, is not enough for Nash play. Aumann and Brandenburger (1995) show that this play requires common knowledge of such rationality, common knowledge of the payoffs, and common knowledge of the conjectured play of the game. Common knowledge of rationality (for example) requires that all agents must know that all agents are rational, and all agents must know that all agents know this, and so on.⁸ In their papers in this issue, Oliver Board and Stephen Morris examine in more detail the connections between game-theoretic solution concepts and the common-knowledge requirements for their use.

This approach is somewhat unsatisfactory. First, common knowledge is a stringent assumption. It might be reasonable to assume that players know their own preferences. Furthermore, it might be reasonable to suppose that they know the preferences of others.

But is it reasonable to suppose that Player A knows that Player B knows that Player A knows that Player B knows Player A's preferences? Perhaps not. Second, the equilibrium selection problem is simply not addressed, since there is common knowledge of conjectured play by assumption. The removal of these common-knowledge requirements and the imposition of only rationality and knowledge of payoffs, however, result only in the prediction of the play of undominated strategies.⁹ Furthermore, common knowledge of rationality and payoffs alone result in the prediction of the play of strategies that survive iterated deletion.¹⁰ Of course, Nash equilibria survive such procedures and hence any selection problem remains.

Alternative assumptions are possible. For example, games of incomplete information formalize the idea that players are perhaps unsure of the payoffs and

rationality of other agents. This approach is embodied in the classic contributions of Harsanyi (1967*a,b*, 1968). Many economic problems fall within this setting. Indeed, the aforementioned examples (such as speculative attacks on currencies and technology choice) are all instances in which the players might not have full knowledge of the payoffs enjoyed by (and hence the motivations of) all the players. In games of incomplete information (when such knowledge is incomplete) behaviour will depend not only on players' beliefs about payoffs, but also on players' beliefs about other players' beliefs, and so on.

Two approaches to this problem have been suggested. First, the problem might be simplified by assuming that the payoffs of each player are drawn independently from a commonly known distribution. Roughly speaking, this means that although the precise payoffs and motivations of an opponent are unknown, the probability that various payoffs arise is commonly known by everyone. In this case, 'higher-order beliefs' (beliefs about the beliefs of others) do not play a role: a player does not know the identity (in the sense of payoffs or, more generally, type) of an opponent, but there is no uncertainty about what beliefs her opponent holds about her. Second, a fully general approach, such as that of Mertens and Zamir (1985), might be taken. This involves a full consideration of a player's beliefs about another player's beliefs about another player's beliefs, and so forth. Unfortunately, this quickly becomes extremely complex and therefore does not lead to sharp answers—and sharper answers are needed to answer policy questions.

In this paper, and in the contributions of this issue, a desire to tackle the problem of multiple equilibria and a dissatisfaction with the imposition of common-knowledge assumptions leads to the study of global games (Carlsson and van Damme, 1993). The specification of global games abandons the common knowledge of payoffs, but does so in a pragmatic way. Formally, Morris and Shin (2002) define global

⁸ For further discussion of common knowledge, see the survey by Geanakoplos (1993).

⁹ A dominated strategy is one that is worse than another strategy, no matter what actions are taken by the other players.

¹⁰ We may delete dominated strategies by discarding any strategies that are dominated. Once we do this, however, it may be the case that other players find that some of their remaining strategies are dominated. The logic proceeds along the following lines: 'My opponent would never find it optimal to play A, therefore I should never play B. She will anticipate this, and hence never play C, so I should never play D.' Such reasoning, when repeated, yields the iterated deletion of dominated strategies. For a standard textbook treatment see, for instance, Osborne and Rubinstein (1994). This procedure has little bite in games with multiple Nash equilibria.

games as: ‘games of incomplete information whose type space is determined by the players each observing a noisy signal of the underlying state’. This definition is somewhat technical, and requires explanation. Formally, there is common knowledge of rationality. In other words, players are assumed to choose optimizing actions given the beliefs they hold about the play of others. However, there is not common knowledge of payoffs—hence we are dealing with ‘games of incomplete information’. Rather, payoffs for all players (which embody their objectives) are determined by some unknown ‘state variable’. A state variable is some aspect of the environment that affects the payoffs of everyone. For instance, in the case of a group of speculative traders attacking a currency, the unknown state variable might be the size of the central bank’s currency reserves. For agents attempting to coordinate on the same technological standard, the state variable might be the true underlying advantages of one technology over another. Although the state variable is unknown, each agent receives a ‘signal’ of it, which provides information about their own payoffs. Crucially, however, it also tells a player about the state variable itself—and hence the likely signals of the other agents. So, if a player were to know the value taken by the state variable, she would understand the ‘real’ game that is in play. Of course, she does not know this, and hence must use her signal (i.e. the information available to her) to infer the value taken by the state variable. Since the signals are not perfectly accurate (so that economic agents are imperfectly informed about the situation), agents may view the game in very different ways. It follows that players must make choices taking into account the fact that other players might not share their views. Thus their strategy choices must be placed in a global context (the class of all games that could possibly be in play) rather than a local context (in which payoffs are commonly known to all players.)

The reasoning employed in the analysis of such global games may appear complex. None the less, the conclusions drawn are often simple. Global games often possess unique equilibria, even though there would be multiple equilibria in a full information analogue: the equilibrium selection problem is (perhaps surprisingly) solved via a relaxation of the common-knowledge assumptions. Furthermore, the

strategies employed in such equilibria correspond closely (and often exactly) to those that a naïve player might choose according to some simple rule of thumb. An explanation of these claims is provided in the remainder of this paper and in the contributions of this issue.

In the next section we consider some simple examples of games with multiple Nash equilibria and the coordination problems that arise. In these examples, the payoffs of the game are assumed to be commonly known by all players.

Following that (section III), we relax this assumption. We allow payoffs to be affected by fundamental uncertainty, in the sense that the payoffs are contingent upon an unobserved state variable. Players privately observe noisy signals of the underlying state variable, upon which they base their actions. This generates strategic uncertainty: players are uncertain of the signals received by others, and hence are uncertain of their likely actions. We use a simple two-action game (interpreted as a technology-adoption problem) to demonstrate that this approach can provide a solution to the equilibrium-selection problem.

We expand this to a many-player example (a public-good contribution problem) in section IV. We explain the roles of fundamental uncertainty (pertaining to the state variable) and strategic uncertainty (pertaining to uncertainty over the actions of others).

Global games require players to infer payoffs from the signals that they receive. Such signals may be privately observed (comprising different signals for each player) or publicly observed (where all players receive the same signal). The nature of equilibrium play depends critically upon the public versus private nature of information sources. This feature is discussed in section V. Finally, section VI provides a brief guide to the themes of this issue.

II. THE PROBLEM OF COORDINATION

The presence of multiple Nash equilibria may be problematic if the play of a particular game is to be predicted. Within the social sciences, equilibrium

multiplicity often manifests itself in the form of a *coordination problem*. In an environment consisting of many interacting decision-making individuals, a Nash equilibrium represents a self-enforcing mode of behaviour. Of course, it is only self-enforcing if individual agents conform to their part in the equilibrium. This might not occur if they worry that other agents will not do the same. In this section, we illustrate this idea with a number of examples drawn from both economics and political science.

(i) Coordination: Technology Adoption

Many products exhibit *network externalities*: their value to a consumer depends (positively) on the level of consumption by others.¹¹ Examples are commonplace. Today, the telephone is a useful device precisely because it has been extensively adopted by others. Electronic mail continues to increase in value as more individuals obtain Internet access.

The different technologies that support new and emerging products often have mutually incompatible standards. In this context, adopting users face a potential coordination problem: it is preferable to adopt the standard that is likely to be adopted by others.

This situation can be represented using a simple 2×2 symmetric game. Suppose that two individuals are faced with the choice of competing standards A and B. Each standard is of use to an agent if and only if the other agent also adopts it. In this case a payoff of $H > 0$ is generated for standard A, or $L > 0$ for standard B, where $L < H$. If the players choose different (incompatible) standards they receive a zero payoff. The strategic-form representation is as follows:¹²

	A	B
A	H	0
B	0	L

By inspection, this game has two pure-strategy Nash equilibria (AA and BB), and adopters face the problem of which standard (equilibrium) to coordinate on. For instance, if both players choose A, then a deviation by a single player will result in the loss of the payoff $H > 0$. In the absence of pre-play communication, societal context, or any other coordination device, the players might very well miscoordinate in their choice of technology, and obtain a zero payoff. Alternatively, societal context might (perhaps for historical reasons) point to the inferior equilibrium. For instance, it is Pareto optimal for players to choose AA, and yet they may coordinate on BB.

One possible argument is that the Pareto-optimal equilibrium AA is *focal*.¹³ Such an argument relies on the fact that it is commonly known that AA is, in fact, Pareto optimal. When players are not commonly certain of the payoffs, however, the row player might worry that the column player believes BB to be Pareto optimal, and hence focal.¹⁴

Examples of such coordination problems have been studied extensively in the economics literature. David (1985) and Liebowitz and Margolis (1990) document the emergence of 'QWERTY' as the dominant configuration for typewriter keyboards.¹⁵ The early 1980s saw a standards battle between the VHS and Betamax formats for domestic video

¹¹ There is a large literature investigating the effect of network externalities. See, for instance, Katz and Shapiro (1985, 1986) and Farrell and Saloner (1985, 1986, 1987), or for a textbook treatment see Shy (1995, 2001). Network externalities are a special case of increasing returns—the marginal benefit from the increased production of a unit increases with the size of the installed base. Increasing returns, and the associated path-dependence of technology choice, have been the subject of much debate. Authors such as Arthur (1989, 1994) and David and Greenstein (1990) have argued that path-dependence may result in lock-in to an inefficient technology. Their arguments are subject to a forceful critique by Liebowitz and Margolis (1994, 1995a).

¹² This is a standard textbook representation. The two players choose an action from the columns and rows, respectively. In each cell (corresponding to a *strategy profile*), the bottom-left payoff is enjoyed by the row player, and the top-right payoff by the column player.

¹³ Focal equilibria are ones which simply 'stand out' from the rest. The idea was originally discussed by Schelling (1960).

¹⁴ The public-good contribution games studied by David Myatt and Chris Wallace in their article in this issue share these features.

¹⁵ David (1985) and Arthur (1989) both cite QWERTY as an example of inefficient lock-in, and extol the virtues of the competing Dworak Simplified Keyboard (DSK). This research has been somewhat discredited by Liebowitz and Margolis (1990), who reveal a number of inconsistencies in the historical sources used by these authors.

recorders.¹⁶ In both of these examples, multiple standards existed for a period of time, but eventually most people coordinated on a single standard. There are currently competing standards in the recordable DVD market, and it is unclear which standard might eventually be adopted to the exclusion of all others. In all of these examples, the number of players is much larger than two. None the less, the simple two-player game above shares many of the properties of a more general multi-player game.

(ii) Many Players: Speculative Currency Attacks

Multi-player coordination games are of great importance at the macroeconomic level. For instance, Cooper (1999) and Cooper and John (1988, 1997) provide a convincing case for the study of coordination problems in macroeconomics. A classic coordination game is provided by the phenomenon of speculative exchange-rate attacks—a recent example of which can be seen in the experience of the United Kingdom in 1992. Simplifying the model of Obstfeld (1996), suppose that a group of speculative traders are considering whether to bet against a currency. Each trader can ‘short sell’ a single unit of currency, at a transaction cost of t . If the currency is devalued, then a trader enjoys a speculative profit of π . If more than a fraction $\hat{\gamma}$ of speculating traders attack (by short selling the currency) then the currency is devalued: the central bank ‘caves in’ in the face of the attack. Writing γ for the actual fraction of speculators choosing to attack the currency, the payoffs for an individual speculator are:

	$\gamma \geq \hat{\gamma}$	$\gamma < \hat{\gamma}$
Attack the currency	$\pi - t$	$-t$
Do not attack	0	0

For different values of the parameters, the game may take different forms. When $\hat{\gamma} > 1$, the currency is never devalued and may be described as ‘stable’. Conversely, for $\hat{\gamma} < 0$, the currency is ‘unstable’. The most interesting region is when $0 < \hat{\gamma} < 1$ and $\pi > t$. In this case, there are multiple pure-strategy

Nash equilibria. It is an equilibrium for all speculators to attack the currency, and also an equilibrium for nobody to attack. Such a region is described by Obstfeld (1996) as one in which the currency is ‘ripe for attack’.

The knowledge the speculators have concerning the values of the parameters is critical. With less than common knowledge it is possible that some agents believe they are playing a game with a single equilibrium, while others believe there are multiple equilibria. Since other agents’ behaviour will depend upon their assessment of the game they are playing, an agent’s beliefs *about their beliefs* will also be critical. Incorporating this insight, Morris and Shin (1998, 1999b) analyse the model of Obstfeld (1996) in a global-game framework (see also Heinemann, 2000). The model is fully discussed and given an experimental treatment by Frank Heinemann in this issue, and a strategically equivalent game is studied in section IV of this Assessment.¹⁷

(iii) Payoff Dominance versus Risk Dominance: Operating System Choice

Personal computers have seen the continued existence of multiple standards. The Microsoft Windows and Apple Macintosh operating systems are both in widespread use. Nowadays, there is extensive compatibility—indeed, this paper was written using both. Nevertheless, there are doubtless advantages to full compatibility, and different groups have coordinated on different operating systems. Most of the business world employs the Microsoft Windows standard, whereas publishing and graphic-design industries make extensive use of Apple systems. Windows is, of course, more widespread than the Apple Mac. For any new computer-system adopters, the coordination problem might take an interesting form. Suppose that two researchers are independently to choose their operating system. Let us further assume that the Apple Mac is technologically superior, so that when both individuals adopt it a higher payoff is enjoyed. In contrast, we suppose that Windows, while inferior, offers a higher security payoff. By this we mean that, in the absence of coordination, Windows affords better opportunities for coopera-

¹⁶ See Liebowitz and Margolis (1995b), who again argue that it is an urban myth that the inferior technology won the standards race.

¹⁷ Other financial settings exhibit similar features. For instance, bank runs (Diamond and Dybvig, 1983) are studied in a global-games framework by Goldstein and Pauzner (2000) and debt pricing is considered by Morris and Shin (1999a).

tion with other computer users. We can represent this idea with the following strategic form game:

	Windows (W)	Mac (M)
Windows (W)	5	2
Mac (M)	4	6

By inspection, we see that there are two pure-strategy Nash equilibria (WW and MM). Notice that in the MM equilibrium each player risks a drop of 4 in his payoff if his colleague fails to coordinate. On the other hand, a drop of only 1 is experienced when failing to coordinate on the WW equilibrium. WW is known as the *risk-dominant* equilibrium (Harsanyi and Selten, 1988). Absent coordination, a tension is observed in this game. Although it might be considered focal to play the (superior) MM equilibrium, a player might also play ‘safe’ by choosing Windows.

To see this a little more formally, suppose that a player has absolutely no idea about the likely actions of his colleague. In fact, the player assigns 50:50 odds to the relative likelihood of his colleague choosing W versus M. Morris and Shin (2002) call such beliefs ‘Laplacian’ following Laplace (1824). Given such beliefs, the expected payoff from Windows is $(5 + 4)/2 = 4.5$. The expected payoff from Mac, however, is $(6 + 2)/2 = 4$. Thus, given these beliefs, it is better to choose Windows rather than Mac. In the context of simple symmetric 2x2 coordination games, this notion of ‘safety’ corresponds exactly to the Harsanyi–Selten (1988) definition of risk dominance. Myatt and Wallace, in this issue, give the more general definition and show how it can be widely applied.

As will be seen, a close correspondence between the risk dominance of a Nash equilibrium and the global-game methodology arises.¹⁸ A game which shares the features of the above operating-system-choice game is analysed later in this issue by Gavin Cameron and Chris Wallace. Here, the risk-dominance

criterion is applied directly to a game with many pure-strategy Nash equilibria in order to select between them. The global-game framework is suppressed from the analysis, but the ideas remain the same.

(iv) Asymmetric Equilibria: Tactical Voting

Coordination problems are not limited to economic settings. In recent parliamentary elections in the United Kingdom, many supporters of the Labour and Liberal Democrat parties wished to ensure the defeat of the Conservative party. In many constituencies, such a defeat required supporters from one of the parties to vote tactically for the other. A stylized version of this situation is one in which the two groups of voters are represented by two different players. For payoffs $H > L > 0$, this situation can be illustrated by the following strategic-form game:

	Labour	Lib Dem
Labour	H	0
Lib Dem	0	L

The top-left outcome corresponds to a Labour win. The bottom-right outcome corresponds to a Liberal Democrat win. Finally, the other outcomes generate a Conservative win following a ‘split’ in the anti-Conservative vote.

Once again there are multiple pure-strategy Nash equilibria. In this case, however, the equilibria are not Pareto ranked. A Liberal Democrat supporter (the row player) would prefer to defeat the Conservative party via a Liberal Democrat win, while a Labour supporter (the column player) would prefer to unseat the Conservatives via a Labour win. Even with common knowledge of payoffs, there would be no simple focal equilibrium argument available. In this issue, David Myatt and Stephen Fisher present a formal model of tactical voting and, by putting the game in a *global* context, obtain a unique equilibrium, and hence derive sharp policy implications.

¹⁸ Morris *et al.* (1995) establish a precise relationship between higher-order beliefs in the presence of less than common knowledge, equilibrium play, and a generalization of risk dominance (*p*-dominance).

The Nash equilibria in the simple coordination games presented here share a common and perhaps unsatisfactory feature: the concept does not yield a unique prediction of play. In the next section, we employ an illustrative example to demonstrate how a global-game framework may be used to provide a solution to the selection problem.

III. THE GLOBAL-GAME APPROACH

Throughout the previous section, it was noted that players might well be uncertain of their environment. Technology adopters may be unsure of the exact benefits generated by each standard. Currency speculators may be unsure of the true floating exchange rate following a devaluation (and hence π), the true state of the government's resolve, or the size of the central bank's currency reserves (both of which will determine γ). Finally, voters in an election may be uncertain of the exact extent of party support.

(i) Modelling Common Uncertainty

Formally, such uncertainty may be modelled as incomplete information over the payoffs of the game. We might proceed in a number of different ways. First, the payoffs of individual players might idiosyncratically vary, but the distribution from which the payoffs are drawn may be commonly known. Formally, this would be a model of 'independent private types'.

A familiar example helps to explain this approach. Independent private value (IPV) auctions are the subject of the celebrated Revenue Equivalence Theorem (Vickrey, 1981; Myerson, 1981). In an IPV auction, a bidder knows her own valuation for an object, and the probabilities of various valuations for opposing bidders. She does not, however, know the exact realization of such opposing valuations. Under such a specification, a player (or bidder) is uncertain of an opponent's payoffs, and hence the action chosen by that opponent. Knowing the distributions of payoffs, however, such a player can deduce the game likely to be played.

In contrast, for richer models of uncertainty, there will be no common knowledge of the game being played. Auctions provide another example. In a (pure) common-value auction, each bidder shares the same (true) valuation for the object question. This common value, however, is unknown. Different bidders may receive different signals of the object's true valuation. This means that the 'types' of the players (a type corresponds to an estimate of the object's value) might be correlated: if one bidder receives a strong signal of the object's value, then it will be more likely that another bidder will also value the object highly.¹⁹

The global-games approach attempts to pursue such 'common value' logic. The basic idea is that the payoffs of the game are determined by some underlying state variable θ . Players then base their decisions on a signal of this underlying state variable. The signal tells them about θ (and hence their own payoffs), but it also tells them about the payoffs of others and, crucially, the signals that others are likely to have received.

(ii) A Global Coordination Game

This idea is best understood via a return to one of the simple coordination games considered in section II. Taking the technology adoption game, and setting $H = L = 1$ yields the following simple pure coordination game:

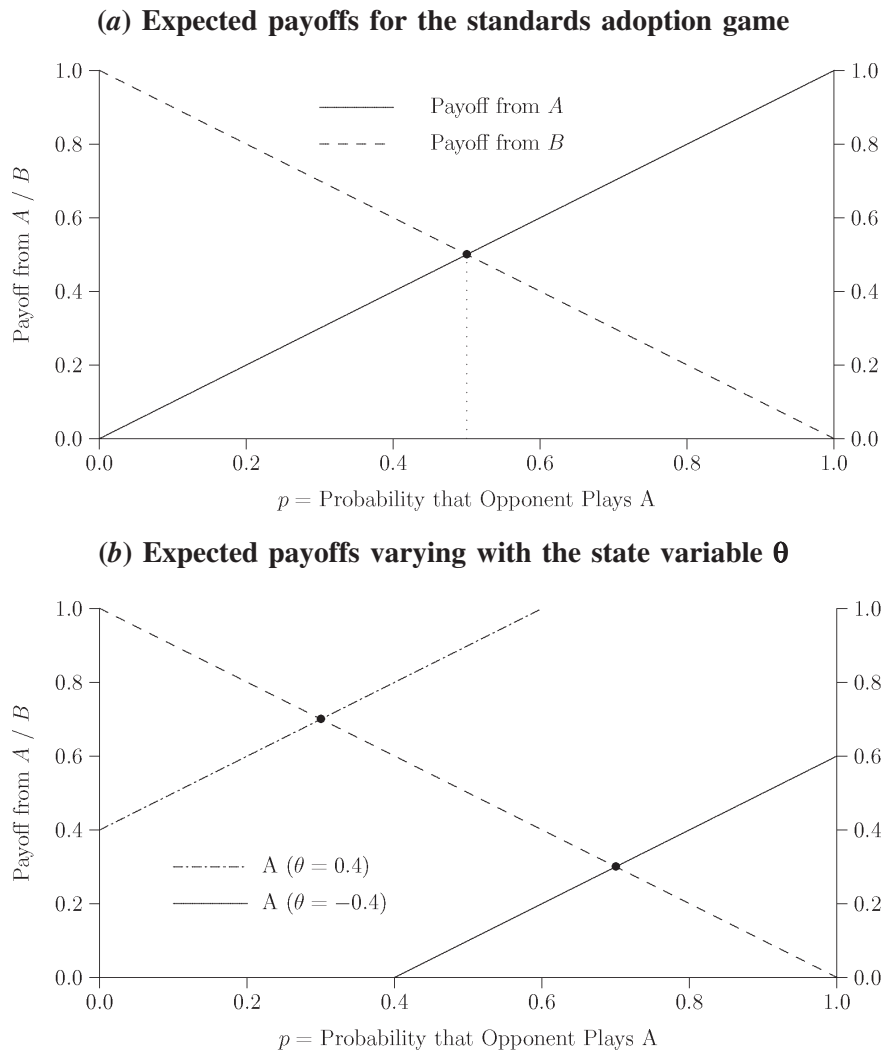
	A	B
A	1	0
B	0	1

A further example of such a coordination game might be the choice of word-processing software. Standards A and B could represent Microsoft Word and WordPerfect respectively.

We have already noted that this game has two pure-strategy Nash equilibria, AA and BB. There

¹⁹ This idea is closely related to the more precise notion of 'affiliation' of players' signals, examined in Milgrom and Weber (1982).

Figure 1
Expected Payoffs in the Technology Adoption Game



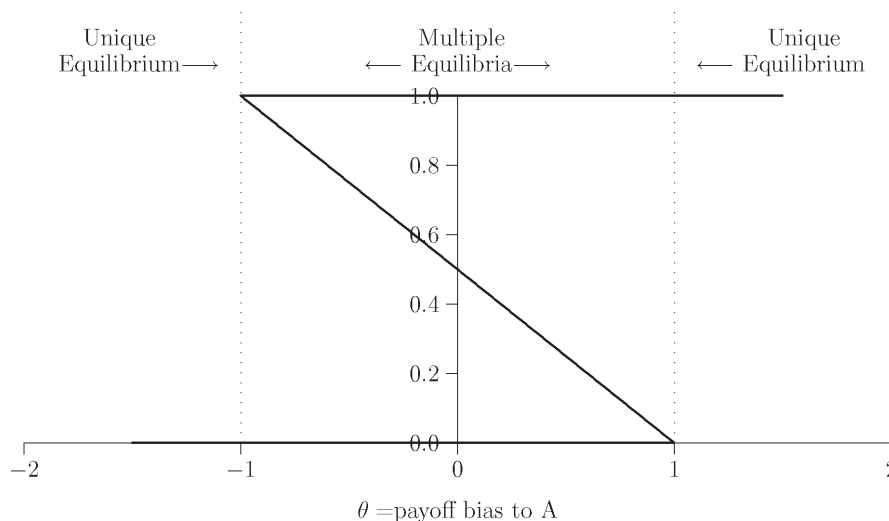
is also a mixed-strategy Nash equilibrium, in which players randomize in their choice of action.²⁰ If a player randomizes 50:50 between the two standards, then her opponent will be exactly indifferent between A and B, and hence willing to randomize in the same way—and this yields an equilibrium.

This idea is illustrated in Figure 1(a). We write p for the probability that an opponent chooses standard A. Thus, the horizontal axis represents the beliefs that a player holds about the likely choice of the other player. The solid line represents the (ex-

pected) payoff that, contingent on these beliefs, the player would receive from choosing standard A. Contrariwise, the dashed line the (expected) payoff, given these same beliefs, that the player would receive from choosing standard B. Hence, given her beliefs, a player should choose the standard that corresponds to the higher of these two lines. Where the solid and dashed lines cross (at $p = 0.5$), a player will be exactly indifferent between the two standards. Thus, such intersections may be used to identify mixed strategy Nash equilibria. The extreme points ($p = 0$ and $p = 1$) correspond to the pure-strategy Nash equilibria.

²⁰ A pure strategy involves a player choosing a fixed action. In a mixed strategy, a player outlines a range of pure strategies (specific actions) and then uses a randomizing device to choose between them. A familiar interpretation (Harsanyi, 1973) is that a player's decision is influenced by an independent payoff shock, so that her decision appears random to her opponent.

Figure 2
Different Games Indexed by θ



Notes: The solid line indicates points that correspond to Nash equilibria for different values of θ . When $\theta \geq -1$, then it is a pure-strategy Nash equilibrium for both players to choose A, corresponding to $p = 1$ on the vertical axis. When $\theta \leq 1$, it is a pure-strategy Nash equilibrium for both players to choose B, corresponding to $p = 0$ on the vertical axis. For $-1 < \theta < 1$, there are mixed-strategy Nash equilibria where players randomize. For instance, when $\theta = 0.5$, it is an equilibrium for players to choose A with probability $p = 0.25$. Doing so ensures that a player is indifferent between A and B, and hence happy to randomize. The mixed-strategy Nash equilibria are illustrated by the downward-sloping portion of the solid line. Notice that for $\theta = 0.5$, it is optimal to play A when a player believes that her opponent will play A with probability greater than 0.25. Thus the downward-sloping line represents a ‘hurdle’ probability. If the opponent is expected to play A with probability greater than this, then it is optimal for the player to join him in this action choice. The dotted lines divide the values of θ into regions with multiple and unique equilibria.

With the basic game in hand, we add a state variable θ . This is a fundamental factor that influences the payoffs of the players. For this example, the state variable will take the form of an additional payoff to standard A, and yields the following strategic-form game:

	A	B
A	$1 + \theta$	0
B	θ	1

In the context of word-processing standard choice, θ might represent the addition of a new feature to standard A (Microsoft Word). For example, θ might represent the additional payoff received owing to

the addition of the ‘Office Assistant’ feature of this program.²¹ If $\theta > 0$, then this feature adds to the usefulness of standard A. In contrast, when $\theta < 0$, it detracts from its usefulness. Finally, when $\theta = 0$, the new feature has no effect, and we return to the original specification. Notice that this payoff shift is common to both players.

According to the value taken by the parameter θ , the game may take a number of different forms. For $\theta > 1$, it is a dominant strategy to choose standard A. For $\theta < -1$, it is a dominant strategy to choose standard B. Finally, when $-1 < \theta < 1$ there are multiple equilibria. As θ moves within this region, the mixed-strategy Nash equilibrium shifts accordingly.

This is illustrated in Figure 1(b). Comparing with Figure 1(a), when $\theta > 0$ the payoff to standard A is higher, shifting the expected payoff line upwards. Similarly, when $\theta < 0$, the payoff to standard A is

²¹ This is better known as the ‘dancing paper clip’ by many Microsoft Word users, and is the subject of some controversy.

lower. The probability p^* associated with the mixed-strategy Nash equilibrium corresponds to the intersection of the payoff lines. By inspection, if $\theta > 0$, then $p^* < 1/2$. This means that, when $\theta > 0$, the best response to 50:50 behaviour on the part of an opponent is to choose standard A. Similarly, when $\theta < 0$, and the opponent's behaviour is completely random, it is better to choose standard B. Allowing θ to range over different values changes the nature of the game, and hence the equilibria. We illustrate this in Figure 2.

If θ is common knowledge (so that everyone knows it, everyone knows that everyone knows it, and so on), then this game may be analysed 'locally'. If it is not, then a player must worry that others will view the game in a different way.

(iii) An Infection Argument

Suppose, for instance, that θ is initially unknown. When a new feature is introduced to a technology standard (such as the addition of the Office Assistant to Microsoft Word) users initially are unaware of its desirability. Suppose, instead, that each player receives information on the likely value of θ . In the case of this example, each player might have the opportunity to inspect standards A and B and evaluate their relative strengths and weaknesses. This allows a player to form an expectation $E[\theta]$, but will not allow her to obtain perfect knowledge of the state variable. In other words, there is *fundamental uncertainty* about the value of the state variable. Nevertheless, a player is able to form expectations about its value.

We now consider a number of different cases. Suppose that a player concludes that $E[\theta] > 1$. Her optimal action is simple: she should choose A. She follows the following logic:

'My signal leads me to estimate $E[\theta] > 1$. Even if my opponent were to play B for sure, it would be optimal to play A. Hence I will play A.'

Suppose instead that $E[\theta] = 0.99$. Inspecting Figure 2, then for fixed $\theta = 0.99$ there would be multiple equilibria. In the absence of common knowledge,

however, a player will worry that her opponent believes that $E[\theta] > 1$. If this is the case, then (when information sources are sufficiently noisy) it will be best for her opponent to choose A. Her logic is as follows:²²

'My signal leads me to estimate $E[\theta] = 0.99$. Hence I should play A if I believe that my opponent will play A with probability 0.01 (1 per cent) or greater. But my signal is imperfect, and it may well be the case that my opponent receives a different signal. It is quite likely (perhaps more than 1 per cent) that my opponent's signal will lead him to conclude that $E[\theta] > 1$, in which case he will play A for sure. So I should play A.'

Notice the presence of *strategic uncertainty*, based upon a lack of common opinions about θ . If our player *knew* that $\theta = 0.99$ (and knew that her opponent knew), multiple equilibria would be present, and there would be no such strategic uncertainty. However, the fact that she does not know that this is the case means that she *suspects* that her opponent may view the game in a way which makes the play of A a dominant strategy. As long as this suspicion is sufficiently large (and by inspection of Figure 2 and the informal argument above, this suspicion need only exceed 1 per cent), then such suspicion is sufficient to ensure the definite play of action A.

This is the first step in an *infection* argument (Morris *et al.*, 1995; Morris and Shin, 1998). Intuitively, the idea is this. A player who believes that $\theta \geq 1$ will choose A. Hence a player who believes that $\theta = 0.99$ will suspect that her opponent believes that $\theta \geq 1$ and hence will also choose A. Now, a player who believes that $\theta = 0.98$ will suspect that her opponent will believe that $\theta \geq 0.99$ and hence will choose A as well:

'My signal leads to me estimate $E[\theta] = 0.98$, and hence I suspect that my opponent might well believe that $E[\theta] \geq 0.99$, in which case he will play A for sure. This is enough to convince me to play A.'

We can continue this argument iteratively. Alternatively, we could begin the argument at $\theta = -0.99$, and consider a range of values for which a player will

²² This argument requires the player to place sufficient probability on the event that her opponent's beliefs are sufficiently different to hers. If this does not hold, then we may simply use the argument with a higher value for $E[\theta]$, say 0.9999.

Box 1
A Formal Model of Signals

Consider a simple formal model of the ‘signal’ structure described in the main text. Suppose that each player is originally ignorant of the parameter θ . This is modelled as a diffuse prior. Player i then receives a signal:

$$\theta_i \sim N(\theta, \sigma^2).$$

Conditional on θ , these signals are independently distributed. Bayesian updating yields posterior beliefs:

$$\theta \mid \theta_i \sim N(\theta_i, \sigma^2) \text{ and } \theta_j \mid \theta_i \sim N(\theta_i, 2\sigma^2).$$

Notice that σ^2 indexes the ‘noise’ in a player’s information sources. For small σ , players have precise beliefs about θ , and for $\sigma \rightarrow 0$ there is almost perfect knowledge of θ . Formally, $1/\sigma^2$ is known as the *precision* of the signal. Using this specification, it is straightforward to observe that a player finds it equally likely that her opponent will have a signal higher or lower than hers. The logic in the text applies, and we will obtain a unique threshold rule with $\theta^* = 0$. For a fixed state variable θ , a player will choose A with probability $\Phi(\theta/\sigma)$, where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal. This means that the two players will successfully coordinate on A with probability $\Phi(\theta/\sigma)^2$. Other coordination and miscoordination probabilities may be calculated similarly, and are illustrated in Figure 3.

find it optimal to play B, working upwards through the different possible configurations of the underlying game.

Such an infection argument leads to the conclusion that a player will play A if and only if $\theta > \theta^*$ for some value θ^* . This is a *threshold rule*, a *cut-off strategy*, or a *switching strategy*. The value θ^* is referred to as the *switching point*. Of course, this is a natural strategy to use: it simply says play action A if and only if there is a sufficiently large payoff bias towards it.

The threshold θ^* used in such a rule will typically be unique. To see why, suppose that a player receives a signal that leads her to the belief that $E[\theta] = \theta^*$. Since this is the ‘threshold’, she should be exactly indifferent between A and B. She reasons:

‘I have a signal leading to the belief that $E[\theta] = \theta^*$. My opponent will choose A if and only if he believes that $E[\theta] \geq \theta^*$. Thus I must ask “What is the probability that his belief is more optimistic than mine?” I have no reason to think that my opponent is more or less optimistic—and hence I will assign probability 0.5 to this event.’


Thus, the player will find it equally likely that her opponent will have beliefs that are more or less

optimistic than hers. To be indifferent between A and B, therefore, it must be the case that $E[\theta] = \theta^* = 0$. In other words, the unique equilibrium threshold rule (and in fact the unique equilibrium) in this ‘global’ version of the standards adoption game is to choose A if and only if θ is perceived to be positive.

Interestingly, this argument reveals a close correspondence between the play of a global game and the risk-dominant (Harsanyi and Selten, 1988) equilibrium of an analogous game with complete information. When $\theta > 0$, the adoption of technology A is a risk-dominant strategy: It is a best response when a player assigns 50:50 odds to the play of her opponent (Figure 2). Such ‘Laplacian’ beliefs might be held by a naïve individual who has no idea about the play of her opponent. Nevertheless, it generates a threshold rule of play that corresponds *exactly* to the *unique* equilibrium way to play this game.

The logic employed here may be applied quite widely. Moreover, in a range of settings the exact structure of players’ signals is not critical (Kajii and Morris, 1997). For the class of 2×2 games (such as the simple coordination game studied here) Carlsson and van Damme (1993) show that, when players receive almost perfect private signals of the payoffs, that play will almost always correspond to the risk-dominant Nash equilibrium of the underlying game.

Figure 3
Probabilities of Coordination and Miscoordination



(a) High precision beliefs: $\sigma = 0.25$

(b) Low precision beliefs: $\sigma = 0.75$

Notes: These figures illustrate the probabilities of observing each possible pure strategy profile (that is, a combination of action choices) in the technology-adoption game. We equip each player with a signal of the state variable θ , generated from the specification described in Box 1. The parameter σ indexes the amount of noise in each signal. When σ is small, the signal is very precise. The argument in section III(iii) demonstrates that a player will choose standard A if and only if her signal is biased towards it. In the figures above, the horizontal axis represents the true value of the state variable θ . For each such θ , we calculate the probability of different signal realizations, and hence action choices, for both players.

Frankel *et al.* (2001) demonstrate that the same logic holds in larger games, particularly those that exhibit strategic complementarities.

(iv) Coordination versus Miscoordination

Our argument establishes that players will successfully coordinate on the same *equilibrium*. Within the context of the example, this does *not* imply that they will coordinate on the same technological standard. The equilibrium strategy specifies a relationship between a player’s beliefs about θ and the action taken. These beliefs are based on the information available to players, and this information may differ. Thus, one player may believe that $\theta > 0$, whereas the other believes that $\theta < 0$, in which case they will choose A and B, respectively, and miscoordinate. This is particularly likely to be true when the true value of θ is close to zero and when the players receive relatively imprecise signals about θ .

We illustrate this idea in Figure 3. Based upon the formal signal specification describe in Box 1, we calculate the probability that players choose each of the different possible action profiles. The ‘noise’ in a player’s information source is indexed by σ . For larger noise, there is a greater probability of miscoordination. Thus, the uniqueness of equilibrium and the model of a player’s information sources yield a valuable analysis of policy: they give us the relationship between the amount of information available about a technological standard and the likelihood that agents will successfully exploit network externalities. Of course, the game in play is directly analogous to the other applications highlighted here and in the other papers in this issue, hence similar intuition will apply in a wider setting. Importantly, however, we have dealt only with a situation in which agents are *privately* informed. The addition of *public* information can, as we show in section V, have a dramatic effect.

IV. FUNDAMENTAL UNCERTAINTY VERSUS STRATEGIC UNCERTAINTY

The discussion so far underscores the importance of the interplay between two types of uncertainty—*fundamental uncertainty* and *strategic uncertainty*. Indeed, this distinction could be regarded as being the organizing principle of all the papers in this issue.

Fundamental uncertainty refers to uncertainty concerning the payoff-relevant state of nature, denoted by θ up to now. Strategic uncertainty refers to the uncertainty concerning the actions of others. This section demonstrates (somewhat formally) that, even as fundamental uncertainty becomes smaller and smaller, strategic uncertainty can remain as large as ever.

(i) A Global Game with Many Players

We can show this in a simple example with a continuum of players (extending the two-player example in section III) who play a public-good contribution game. Each player has to choose between contributing to an indivisible public good, and opting out. For instance, contributing to the public good might be interpreted as paying a membership fee to join a club. All those who join the club enjoy a payoff from a service provided by the club. This service can only be provided if enough people pay their fees.

Let κ be the proportion of players who contribute. The public good is successfully provided when κ is larger than some critical threshold $\hat{\kappa}$. The consumption value of the public good is 1, but player i faces a cost c_i in contributing to the provision of the public good. Thus, the payoffs to player i are as follows:

	$\kappa \geq \hat{\kappa}$	$\kappa < \hat{\kappa}$
Contribute	$1 - c_i$	$-c_i$
Opt out	0	0

Notice the similarity to the currency-speculation game described in section II(ii).²³ The difference is that the contribution cost c_i is indexed by the player i . This is allowed to depend upon a fundamental common component (the state variable) plus noise:

$$c_i = \theta + s_i$$

where θ is the common element in the costs of all players, while s_i is the idiosyncratic ‘payoff shock’ element for player i . The idiosyncratic element s_i is uniformly distributed over the interval $[-\epsilon, \epsilon]$, where ϵ is a small positive number. For any two distinct individuals $i \neq j$, s_i is independent of s_j . Finally, let us suppose that θ itself has a uniform *ex-ante* distribution.

On observing his own cost, player i reasons his way towards the probability density over κ . As a working hypothesis, player i assumes that all other players are using the switching strategy around c^* , so that anyone who has cost below c^* will contribute, while anyone with cost above c^* will opt out. In particular, suppose that player i ’s cost happens to be exactly c^* . Player i then asks himself what the cumulative distribution function over κ is, conditional on c^* . For this, he needs to answer the following question.

‘My cost is c^* . What is the probability that κ is less than z ?’ (Q)

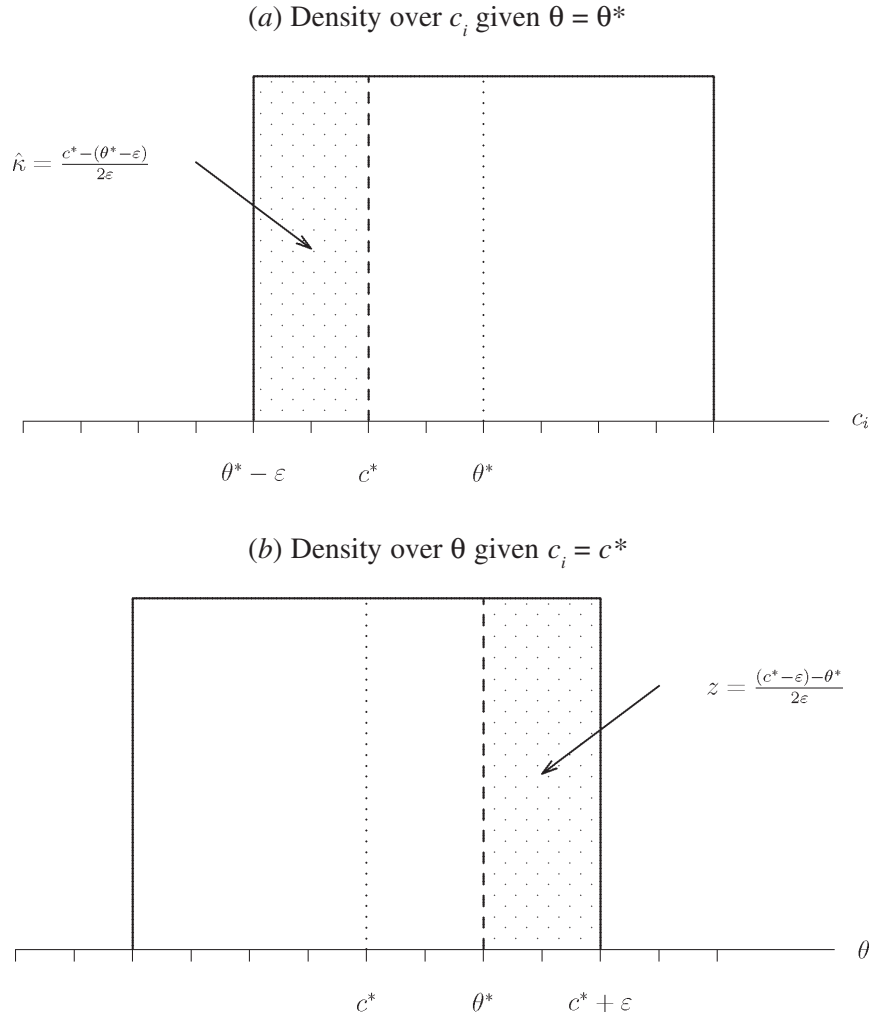
This question is the key to our task, since the answer to question (Q) yields the value of the cumulative

²³ This game is rather different from the public-good provision games studied by Myatt and Wallace in this issue. In the current model, a player must ‘opt in’ in order to benefit from the public good. The specification of Myatt and Wallace supposes that a player enjoys the public good if $\kappa > \hat{\kappa}$, even when she does contribute herself. This would correspond to payoffs:

	$\kappa \geq \hat{\kappa}$	$\kappa < \hat{\kappa}$
Contribute	$1 - c_i$	$-c_i$
Opt out	1	0

The formulation considered here is somewhat easier to analyse, because it exhibits strategic complementarities (the incentive to contribute monotonically increases with κ).

Figure 4
Beliefs about the Fraction of Contributors κ



distribution function of κ evaluated at z . The density over κ is then obtained by differentiating this function. The steps to answering question (Q) are illustrated in Figure 4.

When the common element of cost is θ , the individual costs are distributed uniformly over the interval $[\theta - \epsilon, \theta + \epsilon]$. The players who contribute are those whose costs are below c^* . Hence:

$$\kappa = \frac{c^* - (\theta - \epsilon)}{2\epsilon}.$$

When do we have $\kappa < z$, for a particular choice of z ? This happens when θ is high enough, so that the area under the density to the left of c^* is squeezed out. There is a value of θ at which κ is precisely equal to z . This is when $\theta = \theta^*$, where:


$$\theta^* = c^* + \epsilon - 2\epsilon z.$$

This is illustrated in Figure 4(a). We have $\kappa < z$ if and only if $\theta > \theta^*$. Thus, we can answer question (Q) if we can find the probability that $\theta > \theta^*$.

For this, we must turn to player i 's posterior density over θ conditional on his cost being c^* . This posterior density is uniform over the interval $[c_i - \epsilon, c_i + \epsilon]$. This is because the *ex-ante* distribution over θ is uniform and the idiosyncratic element of cost is uniformly distributed around θ . Figure 4(b) depicts this posterior density over θ . The probability that $\theta > \theta^*$ is then the area under the density to the right of θ^* :

$$\frac{(c^* + \epsilon) - \theta^*}{2\epsilon} = \frac{(c^* + \epsilon) - (c^* + \epsilon - 2\epsilon z)}{2\epsilon} = z.$$

Figure 5
Switching Point c^* as a Function of $\hat{\kappa}$



In other words, the probability that $\kappa < z$ conditional on cost level c^* is exactly z . The cumulative distribution function $G(z | c^*)$ is the identity function:

$$G(z | c^*) = z.$$

The density over κ is then obtained by differentiation:

$$g(\kappa | c^*) = 1 \text{ for all } \kappa.$$

The density over κ is uniform. The noteworthy feature of this result is that the constant ε does not enter into the expression for the density over κ . No matter how small or large is the dispersion of costs, κ has the uniform density over the unit interval $[0, 1]$. Figure 4 reveals the intuition for why ε does not matter. As ε shrinks, the dispersion of costs shrinks with it, but so does the support of the posterior density over θ . The region on the top panel corresponding to z is the mirror image of the region on the bottom panel corresponding to $G(z | c^*)$. Changing ε stretches or squeezes these regions, but does not alter the fact that the two regions are equal in size. This identity is the key to our result.

In the limit as $\varepsilon \rightarrow 0$, every player's cost converges to θ . Thus, fundamental uncertainty disappears. Everyone's cost converges to the common element

θ , and everyone knows this fact. And yet, even as fundamental uncertainty disappears, the strategic uncertainty is unchanged. The subjective density over κ is invariant.

(ii) Switching Points and Strategic Uncertainty

Being able to identify the switching point c^* reveals a lot about the strategic uncertainty that the players face in the game. When $c_i = c^*$, player i is indifferent between contributing and opting out. Denoting by $F(c^*)$ the probability that $\kappa < \hat{\kappa}$ conditional on $c_i = c^*$, the expected payoff to contributing is given by:

$$(1 - c^*)(1 - F(c^*)) - c^*F(c^*) = (1 - c^*) - F(c^*).$$

When player i is indifferent between contributing and opting out, we have:

$$F(c^*) = 1 - c^*.$$

In our case, the density over κ is uniform, so the switching point c^* satisfies:

$$c^* = 1 - \hat{\kappa}$$

Figure 5 illustrates this relationship. The switching point c^* is a decreasing function of the critical mass $\hat{\kappa}$, and takes a particularly simple form.

V. PUBLIC VERSUS PRIVATE INFORMATION

When the informational environment of the economy changes for some reason, both fundamental uncertainty and strategic uncertainty will undergo changes. The net effect on the equilibrium outcome depends on the complex interplay between the two types of uncertainty. This is especially true when the information is *public*, and hence common knowledge among all players. As well as providing information on the underlying state θ , public information also plays the role of conveying information on what others believe and know. By shifting strategic uncertainty in this way, public information may have a big impact on what happens in the game.

We can illustrate this effect informally in the public-good contribution game above. Each player's cost c_i can be seen as a private signal of the common cost component θ . Suppose, however, that the true θ is commonly known to be drawn from a density with mean y . The mean y plays the role of a public signal on θ . It plays a coordinating role in harmonizing the expectations of both players on what values of θ are likely.

Take a simple example. Suppose that y is very high, but player i 's cost is very low. This player then goes through the following reasoning.

'My cost is very low compared to what the mean y of θ would have suggested. There are two possibilities. Either the true realization of θ is very low, and my cost is representative, or the true realization of θ is high, but I have drawn a very low cost. If it is the former (i.e. θ is high), then it is likely that my opponents have drawn high cost levels, which means that they will opt out. However, if it is the latter (i.e. θ is low), then my opponents will have drawn low costs, and choose to contribute. On balance, my signal is very low compared to what I would have expected θ to be. So, I put more weight on θ being high, and my opponents opting out. So, I will choose to opt out.'

So, the mean y of the density from which θ is drawn plays a crucial role. It affects my beliefs about the beliefs held by my opponents, and thereby affects my actions. So, for any cost level drawn by me, I am

more likely to opt out if the mean y is higher. This additional role for the public signal remains even when my own private signal is quite precise.

Similar reasoning applies to more general global games, where the players choose actions that are more finely variable, rather than the simple two-action games mentioned so far. In this issue, Amato, Morris, and Shin examine the impact of public information in a setting where agents have access to public and private information. The agents aim to take actions appropriate to the underlying state, but they also face a spillover effect arising from other agents' actions. Investment decisions, for instance, will be affected by the buoyancy of the market, as well as the underlying fundamentals, and thus will be affected by the actions of other agents. However, social welfare—the sum of the agents' payoff functions—internalizes the spillover effects across all agents. When there is perfect information concerning the underlying state, the unique equilibrium in the game between the agents also maximizes the social welfare function. However, when there is imperfect information, and the agents have access to their own private information, the welfare effects of increased public disclosures turn out to be ambiguous. It is not always the case that greater precision of public information is desirable. In many cases, increased precision of public information is detrimental to welfare. Specifically, the greater the precision of the agents' private information, the greater is the danger posed by increased provision of public information of lowering social welfare.

Amato, Morris, and Shin go on to discuss some of the implications of this result for the conduct of monetary policy. The heightened sensitivity of the market has the potential to magnify any noise in the public information to the degree that public information ends up causing more harm than good. If the information provider anticipates this effect, then the consequence of the heightened sensitivities of the market is to push it into reducing the precision of the public signal. In effect, private and public information end up being *substitutes*, rather than complements.

Among other things, it is not always the case that frequent and timely publication of economic statistics by government agencies and the central bank is

Box 2
The Effect of Public Signals

We can illustrate this more formally by changing the example above so that the common element of cost θ now has a normal distribution with mean y and precision α (i.e. with variance $1/\alpha$). Player i 's cost is given by:

$$c_i = \theta + s_i$$

where the idiosyncratic component of cost s_i is drawn from the normal density with zero mean and precision β . Suppose that all players are following the switching strategy around some point c^* , and that player i 's cost happens to be exactly c^* . We can derive this player's subjective density over κ —the proportion of players who contribute—by following analogous steps to the discussion above. The cumulative distribution function over κ can be obtained from the answer to the following question:

‘My signal is c^* . What is the probability that κ is less than z ?’

The answer to this question will yield $G(z | c^*)$ —the probability that the proportion of players who contribute is at most z , conditional on being at the switching point c^* . Given the common cost element θ , the proportion of players who contribute is:

$$\Phi(\sqrt{\beta}(c^* - \theta))$$

where $\Phi(\cdot)$ is the cumulative distribution function for the standard normal. Let θ^* be the common cost element at which this proportion is exactly z . Thus:

$$\Phi(\sqrt{\beta}(c^* - \theta^*)) = z. \tag{1}$$

When $\theta \geq \theta^*$, the proportion of players that contribute is z or less. So, the question of whether $\kappa \leq z$ boils down to the question of whether $\theta \geq \theta^*$. Conditional on c^* , the density over θ is normal with mean:

$$\frac{\alpha y + \beta c^*}{\alpha + \beta}$$

and precision $\alpha + \beta$. Thus, the probability that $\theta \geq \theta^*$ is the area under this density to the right of θ^* , namely:

$$1 - \Phi\left(\sqrt{\alpha + \beta}\left(\theta^* - \frac{\alpha y + \beta c^*}{\alpha + \beta}\right)\right). \tag{2}$$

This expression gives $G(z | c^*)$. Substituting out θ^* by using equation (1) and rearranging, we can rewrite equation (2) to give:

$$G(z | c^*) = \Phi\left(\frac{\alpha}{\sqrt{\alpha + \beta}}(y - c^*) + \sqrt{\frac{\alpha + \beta}{\beta}}\Phi^{-1}(z)\right).$$

Notice that the mean y of θ now enters the expression, whereas previously it did not. The beliefs over what others do depends critically on what the shared information is.

desirable. By their nature, economic statistics are imperfect measurements of sometimes imprecise concepts, and no government agency or central bank can guarantee flawless information. This raises legitimate concerns about the publication of preliminary or incomplete data, since the benefit arising from early release may be more than outweighed by the disproportionate impact of any error.

The challenge for central banks and other public organizations is to strike the appropriate balance between providing sufficiently accurate signals to the private sector so as to allow it to pursue its goals, but to recognize the inherent limitations of any set of economic statistics and to guard against the potential damage done by the imperfections in the data. This is a difficult balancing act at the best of times.

VI. A GUIDE TO THIS ISSUE

We have noted that game theory has become an integral component of the economist's toolbox. It allows the application of economic thinking to a wider variety of social scientific settings. Unfortunately, many games exhibit multiple equilibria. In addition, conventional interpretations of game theory assume that all relevant aspects of the game in play are commonly known by every player. For game theory to be truly successful, the equilibrium-selection problem must be addressed and the common-knowledge requirements must be abandoned.

The theory of global games offers a pragmatic solution to both of these problems. The common knowledge requirements are relaxed in a tractable manner, via the introduction of *fundamental uncertainty*. Players are then concerned with the fact that other players may view a game in a different way. Such views determine their action choice, and

hence generate *strategic uncertainty*. This forces players to think about all the different games that could be in play—players must think *globally*. Perhaps surprisingly, we demonstrate in this paper that such an approach may yield a *unique* prediction of equilibrium play, that corresponds closely to the play of a naïve player. We obtain, therefore, sharper predictions following the imposition of more realistic modelling assumptions.

This approach permits important policy conclusions. The fact that the equilibrium is unique means that we can perform unambiguous comparative static analyses. For instance, our technology-adoption game allows us to assess the probability with which economic agents will successfully coordinate on the same standard. We also gain an understanding of the relationship between public and private information sources. Importantly, a public information source has a critical coordinating effect. Each player knows that each other player has received the public information, and hence such information can have a dramatic impact upon likely behaviour.

With this in mind, the main theme of this issue is the interplay between uncertainty concerning the underlying fundamentals (the state variable), and uncertainty over the actions of other interested parties (strategic uncertainty). The varied subject matter covered by the papers in this issue bears testimony to the wide-ranging scope of this interplay. We will see in what follows how the formal techniques of global games can be employed to uncover the subtle relationship between the two types of uncertainty—fundamental and strategic—and how the insights can be applied to policy debates that may seem (at first blush) rather far removed from the formalism of games. Needless to say, there are many more avenues to be explored, and we hope this issue and the papers within it can serve to spur more research on the questions addressed here.

REFERENCES

- Arrow, K. J., and Debreu, G. (1954), 'Existence of an Equilibrium for a Competitive Economy', *Econometrica*, **22**, 265–90.
- Hahn, F. (1971), *General Competitive Analysis*, Holden-Day.
- Arthur, W. B. (1989), 'Competing Technologies, Increasing Returns, and Lock-in by Historical Events', *The Economic Journal*, **97**, 642–65.

- Arthur, W. B. (1994), *Increasing Returns and Path Dependence in the Economy*, Ann Arbor, MI, University of Michigan Press.
- Aumann, R., and Brandenburger, A. M. (1995), 'Epistemic Conditions for Nash Equilibrium', *Econometrica*, **63**(4), 1161–80.
- Bertrand, J. (1883), 'Review of *Théorie Mathématique de la Richesse Sociale* by Léon Walras; and of *Recherches sur Les Principes Mathématiques de la Théorie Des Richesses* by Augustin Cournot', *Journal des Savants*, **67**, 499–508.
- Carlsson, H., and van Damme, E. (1993), 'Global Games and Equilibrium Selection', *Econometrica*, **61**(5), 989–1018.
- Cooper, R. W. (1999), *Coordination Games: Complementarities and Macroeconomics*, Cambridge, Cambridge University Press.
- John, A. (1988), 'Coordinating Coordination Failures in Keynesian Models', *Quarterly Journal of Economics*, **103**, 441–63.
- — (1997), 'Dynamic Complementarities: A Quantitative Analysis', *Journal of Monetary Economics*, **40**, 97–119.
- Cournot, A. (1838), *Recherches sur Les Principes Mathématiques de la Théorie Des Richesses*, New York, Macmillan, as *Researches into the Mathematical Principles of the Theory of Wealth*, English edn (1897).
- David, P. A. (1985), 'Clio and the Economics of QWERTY', *American Economic Review*, **75**, 332–7.
- Greenstein, S. (1990), 'The Economics of Compatibility Standards: An Introduction to Recent Research', *Economics of Innovation and New Technology*, **1**, 3–42.
- Debreu, G. (1959), *Theory of Value: An Axiomatic Analysis of Economics Equilibrium*, New Haven, CT, and London, Yale University Press.
- Diamond, P., and Dybvig, P. (1983), 'Bank Runs, Deposit Insurance, and Liquidity', *Journal of Political Economy*, **91**, 401–19.
- Edgeworth, F. Y. (1897), 'The Pure Theory of Monopoly', in F. Y. Edgeworth (ed.), *Papers Relating to Political Economy*, London, Macmillan (1925).
- Farrell, J., and Saloner, G. (1985), 'Standardization, Compatibility, and Innovation', *RAND Journal of Economics*, **16**, 70–84.
- — (1986), 'Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation', *American Economic Review*, **76**, 940–55.
- — (1987), 'The Economics of Horses, Penguins, and Lemmings', in L. Gable (ed.), *Production Standardization and Competitive Strategies*, Amsterdam, North-Holland.
- Frankel, D. M., Morris, S., and Pauzner, A. (2001), 'Equilibrium Selection in Global Games with Strategic Complementarities', *Journal of Economic Theory*, forthcoming.
- Fudenberg, D., and Levine, D. (1998), *The Theory of Learning in Games*, Cambridge, MA, MIT Press.
- Geanakoplos, J. (1993), 'Common Knowledge', ch. 40 in R. Aumann and S. Hart (eds), *Handbook of Game Theory*, Amsterdam, Elsevier Science.
- Goldstein, I., and Pauzner, A. (2000), 'Demand Deposit Contracts and the Probability of Bank Runs', available from <http://www.tau.ac.il/pauzner>
- Harsanyi, J. C. (1967a), 'Games with Incomplete Information Played by "Bayesian" Players: Part I', *Management Science*, **13**, 159–82.
- (1967b), 'Games with Incomplete Information Played by "Bayesian" Players: Part II', *Management Science*, **13**, 320–34.
- (1968), 'Games with Incomplete Information Played by "Bayesian" Players: Part III', *Management Science*, **14**, 486–502.
- (1973), 'Games with Randomly Disturbed Pay-offs: A New Rationale for Mixed-strategy Equilibrium Points', *International Journal of Game Theory*, **2**, 1–23.
- Selten, R. (1988), *A General Theory of Equilibrium Selection in Games*, Cambridge, MA, MIT Press.
- Heinemann, F. (2000), 'Unique Equilibrium in a Model of Self-fulfilling Currency Attacks: Comment', *American Economic Review*, **90**, 316–18.
- Hofbauer, J., and Sigmund, K. (1988), *The Theory of Evolution and Dynamical Systems*, London, Cambridge University Press.
- Kajii, A., and Morris, S. (1997), 'The Robustness of Equilibria to Incomplete Information', *Econometrica*, **65**, 1283–309.
- Katz, M., and Shapiro, C. (1985), 'Network Externalities: Competition and Compatibility', *American Economic Review*, **75**, 424–40.

- Katz, M., and Shapiro, C. (1986), 'Technology Adoption in the Presence of Network Externalities', *Journal of Political Economy*, **94**, 822–41.
- Kreps, D. M. (1990), *A Course in Microeconomic Theory*, New York, Harvester Wheatsheaf.
- Kuhn, H. W. (ed.) (1997), *Classics in Game Theory*, Princeton, NJ, Princeton University Press.
- Laplace, P. (1824), *Essai Philosophique sur les Probabilités*, Dover (English translation).
- Liebowitz, S., and Margolis, S. E. (1990), 'The Fable of the Keys', *Journal of Law and Economics*, **33**(1), 1–25.
- (1994), 'Network Externality: An Uncommon Tragedy', *Journal of Economic Perspectives*, **8**(2), 133–50.
- (1995a), 'Are Network Externalities a New Source of Market Failure?', *Research in Law and Economics*, **17**, 1–22.
- (1995b), 'Path Dependence, Lock-in, and History', *Journal of Law, Economics and Organization*, **11**(1), 205–26.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995), *Microeconomic Theory*, Oxford and New York, Oxford University Press.
- Maynard Smith, J. (1974), 'The Theory of Games and the Evolution of Animal Conflicts', *Journal of Theoretical Biology*, **47**, 209.
- Mertens, J.-F., and Zamir, S. (1985), 'Formulation of Bayesian Analysis for Games with Incomplete Information', *International Journal of Game Theory*, **10**, 619–32.
- Milgrom, P. R., and Weber, R. (1982), 'A Theory of Auctions and Competitive Bidding', *Econometrica*, **50**(5), 1089–122.
- Morris, S., and Shin, H. S. (1998), 'Unique Equilibrium in a Model of Self-fulfilling Currency Attacks', *American Economic Review*, **88**(3), 587–97.
- (1999a), 'Coordination Risk and the Price of Debt' *European Economic Review*, forthcoming, available at <http://www.econ.yale.edu/smorris>
- (1999b), 'A Theory of the Onset of Currency Attacks', in P.-R. Agenor, M. Miller, D. Vines, and A. Weber (eds), *Asian Financial Crises: Causes, Contagion and Consequences*, Cambridge, Cambridge University Press.
- (2002), 'Global Games: Theory and Application', in M. Dewatripont, L. P. Hansen, and S. J. Turnovsky (eds), *Advances in Economics and Econometrics: Theory and Applications*, London, Cambridge University Press.
- Rob, R., and Shin, H. S. (1995), ' p -Dominance and Belief Potential', *Econometrica*, **63**(1), 145–57.
- Myerson, R. (1981), 'Optimal Auction Design', *Mathematics of Operations Research*, **6**(1), 58–73.
- (1999), 'Nash Equilibrium and the History of Economic Theory', *Journal of Economic Literature*, **37**, 1067–82.
- (2002), 'History of Nash Equilibrium', inaugural lecture at the Social Sciences Division of the University of Chicago.
- Nash, J. F. (1950a), 'Equilibrium Points in N -person Games', *Proceedings of the National Academy of Sciences*, **36**, 48–9.
- (1950b), 'Non-cooperative Games', Ph.D. thesis, Princeton University.
- (1951), 'Non-cooperative Games', *Annals of Mathematics*, **54**(2).
- Obstfeld, M. (1996), 'Models of Currency Crises with Self-Fulfilling Features', *European Economic Review*, **40**, 1037–47.
- Osborne, M. J., and Rubinstein, A. (1994), *A Course in Game Theory*, Cambridge, MA, MIT Press.
- Robbins, L. (1932), *Essay on the Nature and Significance of Economic Science*, London, Macmillan.
- Samuelson, L. (1997), *Evolutionary Games and Equilibrium Selection*, Cambridge, MA, MIT Press.
- Schelling, T. C. (1960), *The Strategy of Conflict*, Cambridge, MA and London, Harvard University Press.
- Shy, O. (1995), *Industrial Organization*, Cambridge, MA, MIT Press.
- (2001), *The Economics of Network Industries*, Cambridge, Cambridge University Press.
- Vega-Redondo, F. (1996), *Evolution, Games, and Economic Behaviour*, Oxford, Oxford University Press.
- Vickrey, W. (1981), 'Counterspeculation, Auctions, and Competitive Sealed Tenders', *Journal of Finance*, **16**(1), 8–37.
- von Neumann, J., and Morgenstern, O. (1967), *Theory of Games and Economic Behaviour*, 3rd edn, New York, Wiley.
- Weibull, J. W. (1995), *Evolutionary Game Theory*, Cambridge, MA, MIT Press.
- Young, H. P. (1998), *Individual Strategy and Social Structure*, Princeton, NJ, Princeton University Press.