3 COORDINATION PROBLEMS

In the previous section, we focused on individual behavior in dynamic intertemporal optimization problems where the optimal rational expectations solution was unique. In many macroeconomic environments, this is not the case. Instead, multiple rational expectations equilibria exist, and the question is which of these equilibria economic agents will choose to coordinate upon. Laboratory experiments can be quite useful in this regard. Indeed, Lucas (1986) argued that laboratory experiments were a reasonable means of resolving such coordination problems because "economic theory does not resolve the situation [so] it is hard to see what can advance the discussion short of assembling a collection of people, putting them in the situation of interest, and seeing what they do."

Some coordination problems of interest to macroeconomists were previously addressed in Ochs (1995). In particular, that chapter surveyed experimental studies of overlapping-generations models, where money may or may not serve as a store of value (Lim, Prescott, and Sunder 1994) or subjects can select between low- or high-inflation equilibria (Marimon and Sunder 1993, 1994, 1995). Also included were experimental studies of stag-hunt and battle-of-the sexes games (surveyed also in Cooper (1999) and Bryant (1983)–type Keynesian coordination games (e.g., the minimum- and median-effort games of Van Huyck , Battalio, and Beil (1990, 1991) and Van Huyck, Cook, and Battalio (1994).[20] The coordination games literature delivered a number of important findings on when coordination success was likely to be achieved and when coordination failure was likely. Importantly, the results have been replicated by many other experimenters, leading to confidence in those findings. Rather than review those replications and extensions, in this section I report on more recent macrocoordination experiments. The environments tested in these experiments have a more direct resemblance to macroeconomic models than do the coordination games surveyed by Ochs (with the exception of Marimon and Sunder's work on overlapping generations models). I also address some equilibrium-selection mechanisms or refinements that have been proposed for resolving macrocoordination problems and the experimental studies of those mechanisms and refinements.

3.1 Poverty Traps

Lei and Noussair (2007) build on their (2002) experimental design for studying behavior in the one-sector optimal growth model by adding a nonconvexity to the production technology, resulting in multiple, Pareto-rankable equilibria. Specifically, the production function used to determine output in Noussair and Matheny (2000) and Lei and Noussair (2002) is changed to

$$f(k_t) = \begin{cases} \underline{A}k_t^{\alpha} & \text{if } k_t < k^* \\ \overline{A}k_t^{\alpha} & \text{if } k_t \geq k^* \end{cases}$$

where $\underline{A} < \overline{A}$ and k^* is a threshold level of the aggregate capital stock that is known to all five subjects. The threshold switch in productivity is a simple way of modeling positive externalities that may arise once an economy reaches a certain stock of capital (physical or human; see, e.g., Azariadis and Drazen 1990). An implication is that there are now two stationary levels for the capital stock (and output) $\overline{k}_l < k^* < \overline{k}_h$, with $\overline{k}_l$

representing the poverty trap and $\bar{k}_h$ representing the Pareto efficient equilibrium. The dynamics of the system (under perfect foresight) are such that for $k \in (0, k^*)$, $\bar{k}_l$ is an attractor, whereas for $k \geq k^*$, $\bar{k}_h$ is the attractor. The main experimental question is on which of these two equilibria subjects will learn to coordinate.

One treatment variable was the initial aggregate level of the capital stock, either below or above the threshold level k^* and divided up equally among the five subjects. The other treatment condition was whether decisions were made in a decentralized fashion, with a market for the capital stock (subjects had different production technologies that aggregated up to the aggregate technology), or whether groups of subjects together made a collective consumption-savings decision, that is, playing the role of a social planner. In both cases, the indefinite horizon of the model was implemented using a constant probability of continuation, and subjects were paid on the basis of the utility value of the consumption they were able to achieve in each period. The main experimental finding is that in the decentralized treatment, the poverty-trap equilibrium is a powerful attractor; it is selected in all sessions where the initial aggregate capital stock is below k^* as well as in some sessions where the initial aggregate capital stock lies above k^*. There are some instances of convergence to the Pareto efficient stationary equilibrium $\bar{k}_h$ but only in the decentralized setting, where the initial capital stock lies above k^*. In the social planner treatment, where five-subject groups jointly decide on consumption-savings decisions, neither of the two stationary equilibria were ever achieved; instead there was either convergence to a capital stock close to the threshold level k^* or to the golden-rule level that maximally equates consumption in every period. While the latter is close to the Pareto optimum, it is inefficient as it ignores the possibility that the economy may terminate (the rate of time preference is positive). Lei and Noussair (2007) conclude that additional institutional features may be necessary to both avoid and escape from the poverty-trap outcome.

The possibility that various institutional mechanisms might enable economies to escape poverty traps is taken up in a follow-up experimental study by Capra and others (2009). These authors begin by noting that laboratory studies of the role of institutions in economic growth may avoid endogeneity problems encountered in field data studies (where it is unclear whether institutions cause growth, or vice versa) and more clearly explore environments with multiple institutions. The two institutions explored in this study are termed *freedom of expression,* which involves free discussion among subjects prior to each round of decision making and *democratic voting,* in which subjects vote on two proposals for how to divide output up between consumption and savings (future capital) at the end of each period.

The baseline experimental design is essentially the same as the low initial capital stock treatment of Lei and Noussair (2007); there are five subjects who begin each indefinite sequence of rounds with capital stocks that sum up to an aggregate level that lies below the threshold level k^*.[21] This initial condition for the aggregate capital stock is the same in all treatments of this study because the focus here is on whether subjects can escape from the poverty-trap equilibrium. At the start of a period, output is produced based on last period's capital stock, and then a market for capital (the output good) opens. After the market for capital has closed, subjects independently and without communication decide on how to allocate their output between current consumption and savings (next period's capital stock). In the communication treatment, subjects are free to communicate with one another prior to the opening of the market for capital. In the voting treatment, after the capital market has closed, two subjects are randomly selected to propose consumption/savings plans for all five agents in the economy; these

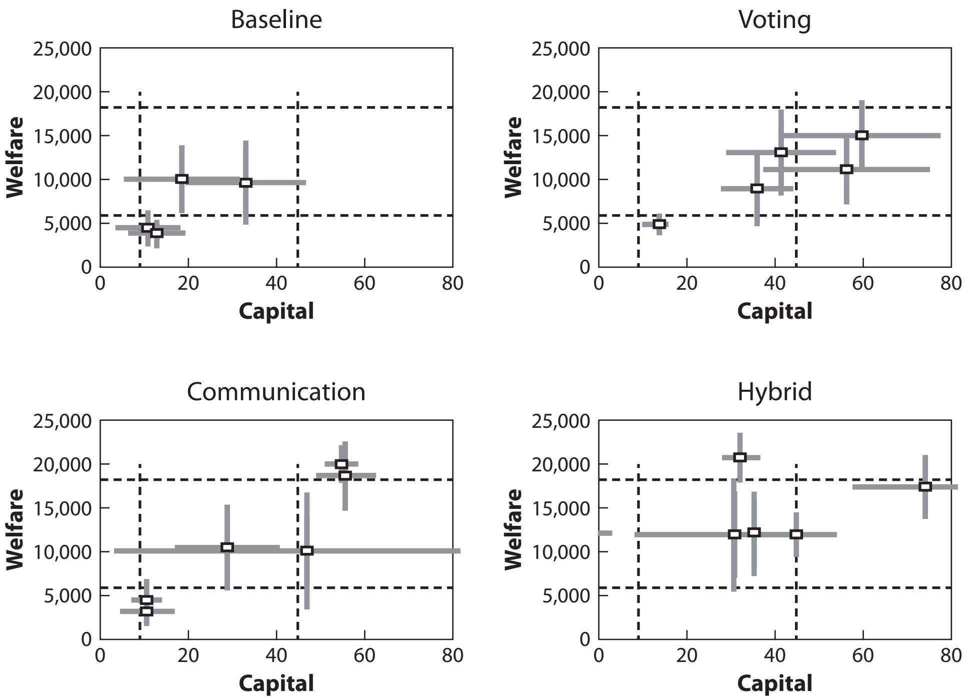

Figure 1.4: Asymptotic estimates of aggregate welfare (vertical axis) and capital (horizontal axis) for each session (square) of the four treatments of Capra et al. (2009).

proposals specify how much each subject is to consume and how much to invest in next period's capital stock (if there is a next period). Then all five subjects vote on the proposal they prefer, and the proposal winning a majority of votes is implemented. In a hybrid treatment, both communication and voting stages are included together.

The main findings examine the long-run values of two statistics for each session: (1) aggregate welfare (as measured by the sum of the period utility from consumption by all five agents $\sum_i u(c_t^i)$) and (2) the aggregate capital stock ($\sum_i k_t^i$). Capra and others (2009) use an equation similar to (1) to estimate the asymptotic values of these two measures for each five-person economy.[22] These estimated values are shown as squares in Figure 1.4, and the line segment through each square represents the 95% confidence region. The lower-left intersection of the dashed lines shows the poverty-trap level of aggregate welfare and capital, while the upper-right intersection of the two dashed lines shows the Pareto efficient level of aggregate welfare and capital. This figure reveals the main findings in the baseline treatment: consistent with Lei and Noussair, subjects are unable to escape from the poverty-trap outcome. The addition of communication or voting helped some, though not all, economies to escape from the poverty trap. In the hybrid model, which allows both communication and voting, the experimental economies appear to always escape from the poverty trap (95% confidence bounds exclude poverty-trap levels), and these economies are closest to the Pareto efficient equilibrium levels for welfare and the capital stock. Capra and others argue that binding consumption/savings plans as in the voting treatment are important for achieving aggregate capital stock levels in excess of the threshold level, while communication makes it more likely that such consumption/savings plans are considered in the first place; not surprisingly then, the two institutions complement one another well and lead to the best outcomes.

While this experimental design involves a highly stylized view of the institutions labeled "freedom of expression" and "democratic voting," the same critique can be made of the neoclassical model of economic growth. The experimental findings suggest that there may be some causality from the existence of these institutions to the achievement of higher levels of capital and welfare, though the opposite direction of causality from growth to institutions remains an important possibility. More recently, macroeconomists have emphasized the role of human capital accumulation, so it would be of interest to consider whether subjects learn to exploit a positive externality from a highly educated workforce. And while several other studies have pointed to the usefulness of communication in overcoming coordination problems (e.g., Blume and Ortmann 2007; Cooper et al. 1992), these have been in the context of strategic form games. While the results of those studies are often cleaner, in the sense that the game is simple and communication is highly scripted, the study by Capra and others implements institutional features in a model that macroeconomists care about, and this may serve to improve the nascent dialogue between experimentalists and macroeconomists.

3.2 Bank Runs

Another important coordination problem that has been studied experimentally in the context of a model that macroeconomists care about is Diamond and Dybvig's (1983) coordination game model of bank runs. In this three-period intertemporal model, depositors find it optimal to deposit their unit endowment in a bank in period 0, given the bank's exclusive access to a long-term investment opportunity and the deposit contract the bank offers. This deposit contract provides depositors with insurance against uncertain liquidity shocks; in period 1, some fraction of depositors learn they have immediate liquidity needs (are impatient) and must withdraw their deposit early, while the remaining fraction learn they are patient and can wait to withdraw their deposit in the final period 2. The bank uses its knowledge of these fractions in optimally deriving the deposit contract, which stipulates that depositors may withdraw the whole of their unit endowment at date 1, while those who wait to withdraw until period 2 can earn $R > 1$. While there exists a separating, Pareto efficient equilibrium where impatient types withdraw early and patient types wait until the final period, there also exists an inefficient pooling equilibrium, where uncertainty about the behavior of other patient types causes all patient types to mimic the impatient types and withdraw their deposits in period 1 rather than waiting until period 2. In the latter case, the bank has to liquidate its long-term investment in period 1; depending on the liquidation value of this investment, it may have insufficient funds to honor its deposit contract in period 1. The possibility of this bank-run equilibrium is the focus of experimental studies by Garratt and Keister (2009), Schotter and Yorulmazer (2009), Madiés (2006), and Arifovic, Jiang, and Xu (2013). All of these experiments dispense with inducing the two player types and focus on the decisions of the single "patient" player type alone, who is free to choose whether to run on the bank (mimicking an impatient type) or not, that is, they all focus on the pure coordination game aspect of the problem.

Garratt and Keister (2009) study the coordination game played by five subjects who have \$1 deposited in a bank and must decide at one or more opportunities whether to withdraw their \$1 or leave it deposited in the bank, potentially earning a higher return of \$1.50. Following each withdrawal opportunity, subjects learn the number of players

TABLE 1.2:
Bank-run coordination game payoffs. Source: Garratt and Keister (2009).

Hypothetical No. of Withdrawl Requests	*Amount Each Requester Would Receive*	*Projected Payment to Each Depositor*
0	n/a	$1.50
1	$1	$1.50
2	$1	$1.50
3	$1	$0
4	$0.75	$0
5	$0.60	n/a

in their group of five (if any) who have chosen to withdraw. As treatment variables, Garratt and Keister varied the number of withdrawal opportunities (one or three) and the number of early withdrawals a bank could sustain while continuing to offer those who avoided withdrawal a payoff of $1.50 (i.e., variation in the liquidation value of the bank's long-term investment). Table 1.2 provides one parameterizations of Garratt and Keister's bank-run game.

Garratt and Keister report that for this baseline game, regardless of the liquidation value of the long-term investment, no group ever coordinated on the "panic equilibrium" (five withdrawals) and a majority of groups coordinated on the payoff dominant equilibrium (zero withdrawals). In a second treatment that more closely implements the liquidity shock in the Diamond-Dybvig model, Garratt and Keister added "forced withdrawals" to the baseline game: at each withdrawal opportunity, there was a small known probability that one randomly selected player would be forced to withdraw. However, whether a withdrawal was forced or not was unknown to subjects. The probabilities of forced withdrawals were chosen such that there continued to exist a payoff dominant equilibrium in which no player ever voluntarily withdrew at any withdrawal opportunity (if all adhered to this strategy, they would earn an expected payoff greater than $1) as well as a panic equilibrium where all withdraw. Garratt and Keister report that with forced withdrawals (liquidity shocks), the frequency of voluntary withdrawals and coordination on the panic equilibrium is significantly greater relative to the baseline treatment with unforced withdrawals. This increase in panic behavior was particularly pronounced in the forced-withdrawal treatment where subjects had multiple withdrawal opportunities and could condition their decisions on the prior decisions of others. An implication of this finding is that panic behavior may require some conditioning on the decisions of others, suggesting that the bank-run phenomenon is perhaps best modeled as a dynamic game, as opposed to the simultaneous-move formulation of Diamond and Dybvig (1983).

Schotter and Yorulmazer (2009) arrive at a similar conclusion, using a somewhat different experimental design. Theirs involves a group of six subjects deciding in which of four periods to withdraw their deposit of $\$K$ in the face of uncertainty concerning both the withdrawal decisions of the other five subjects as well as the type of bank in which all six have invested their deposits. Subjects know that there are five possible bank types, that each type is equally likely to be drawn for the duration of each four-period game, and that the mean return across types is r^*.[23] While the bank type is unobservable, the "promised" return is fixed at 12% per period, while the mean return r^* was varied across sessions, either 0.07, 0.08 or 0.14. Subjects were told that if they kept their $\$K$ deposit invested for ℓ periods, they could earn a return of $\$(1.12)^{\ell}K$ if the bank has sufficient funds left in period ℓ, but if not, the bank would pay all those withdrawing in that period an equal share of remaining funds on hand (if any). Subjects had to choose in which of the four periods to withdraw their money, with withdrawal being irreversible. The authors think of this as a model of a bank run in progress (the precipitating event is left unmodeled) and are interested in exploring three factors that may slow or hasten the period in which deposits are withdrawn. A first factor is whether the withdrawal decision across the four periods is implemented as a simultaneous-move normal-form game or as an extensive-form game; in the former case subjects specify the period in which they want to withdraw their funds (1, 2, 3, or 4), while in the latter case subjects make withdrawal decisions period by period and may condition on the prior period withdrawal decisions (and in one treatment, the amounts earned) by others. The second and third factors are the use of deposit insurance to delay or slow down the run or the presence of insiders who know the mean return r^* of the banks and may, through their actions, persuade other uniformed subjects to run early or wait.

Schotter and Yorulmazer (2009) find that bank runs are less likely to be severe (withdrawal occurs later, e.g. in period 3 or 4) when r^* is known to be greater than the bank's promised return of 12%. For fixed r^*, runs are also less severe in the extensive-form version of their model, when agents can condition on the decisions of others and there is a high degree of information, in that subjects also know the amounts that others have received.[24] This finding is interesting in that theory does not predict that the game form should matter; the fact that it does again points to the value of thinking of bank runs as dynamic rather than static games. They further show that partial-deposit insurance may work to diminish the severity of bank runs, as can the presence of some depositor insiders who know the type of bank with which funds have been invested.

Madiés (2006) examines bank runs as two-period pure coordination games repeatedly played (thirty repetitions) by larger groups of ten subjects. Madiés varied (1) the difference in payoffs from early versus late withdrawals, (2) the number of early withdrawals a bank could sustain while continuing to offer those who avoided an early withdrawal their promised late-withdrawal payment, and (3) the role played by suspension of deposit availability (implemented as suspension of activity during the experiment to calm the panic) or deposit insurance of either 25% or 75% coverage in arresting bank runs. Among other findings he reports that pure panic equilibria, where all ten subjects run in the first period, are rare under all treatment conditions and that partial runs are much more common, even though such partial runs are not equilibria of the model. Further, threatened suspensions of deposit availability are rather effective at preventing bank runs, while partial-deposit insurance is essentially ineffective.

Arifovic, Jiang, and Xu (2013) also study two-period bank runs as pure coordination games with groups of ten subjects. They fix the pure strategy run equilibrium payoff to 1 and the pure strategy no-run equilibrium payoff to 2 and systematically vary the short-run return to early withdrawal, which can be reinterpreted as a coordination parameter, η, specifying the minimum fraction of depositors who must withdraw late to equalize the payoffs earned from early and late withdrawals. Their main finding is that runs reliably occur when η is 0.7 or greater, that is, when at least 70% of subjects must withdraw late in order to achieve a payoff that is at least as high as the payoff from withdrawing early. One novelty of their design is that they do not use neutral language and frame the game played as a decision of when to withdraw deposits from a bank.

The issue of the *contagious spread* of a bank run from one location to another is addressed experimentally by Corbae and Duffy (2008). They study a two-stage, four-player game. In the first stage, players simultaneously propose to form links with one another; mutually agreeable links are then implemented and comprise the set of each player's "neighbors." Corbae and Duffy interpret the players as "banks" connected to one another via interbank reserve deposits that can serve to insure against risk. (à la Allen and Gale 2000). In the second stage, each player plays τ rounds of an n-person, equal-weighted-payoff "stag"-hunt game with his $n = 1, 2,$ or 3 neighbors. As in Garratt and Keister (2009), one of the four-players is "shocked"—that is, randomly must play the inefficient "hare" or run strategy in all rounds of the second-stage game. Corbae and Duffy define a contagion as a movement by all players away from the Pareto efficient stag equilibrium to the inefficient hare equilibrium. While it is possible for subjects to implement a complete network of links (the four players have three links each) that provides insurance against the risk of being linked to a player forced to panic, as when all unshocked players play stag, Corbae and Duffy show that such a network configuration is not an equilibrium due to the free-rider problem. Instead, the network configurations that are predicted to emerge are *bilateral* networks (two-player networks, where each player has a single link), which serves to limit the spread of the bank-run outcome. Corbae and Duffy report experimental evidence that is broadly consistent with this prediction. Starting groups of four subjects out in different exogenous network configurations and then in subsequent games allowing them to choose the players they want to link to, they report that subjects consistently move in the direction of choosing to have a single link to one other player. Under this bilateral network, the bank-run equilibrium is isolated to just one of the two-player networks; the other network achieves the efficient, payoff-dominant equilibrium.

Summing up, we have discussed two kinds of macroeconomic-coordination experiments, poverty traps and bank runs. In the poverty-trap model, the question of interest is how to get subjects to move *from* an inefficient equilibrium *to* an efficient one. We might think of this as a good contagion. In the bank-run model, the question of interest is precisely the opposite: how to keep funds deposited in a bank longer (earning higher returns) and avoid a bad contagion to an inefficient panic equilibrium. Both types of movements are difficult to achieve in the laboratory. In the case of movement from an efficient to an inefficient equilibrium, it seems necessary to force some players' hands in order to precipitate a transition to the inefficient outcome; that finding suggests that the precise mechanism precipitating a bad contagion has yet to be discovered.

We next explore experimental tests of two mechanisms that macroeconomists have used to resolve coordination problems.

3.3 *Resolving Coordination Problems: Sunspots*

In the bank-run coordination game, the question of equilibrium selection is left unmodeled. Diamond and Dybvig (1983) suggest that depositors might use realizations of some commonly observed, nonfundamental random variable, or "sunspot" in the language of Cass and Shell (1983) and Azariadis (1981), to resolve the question of the equilibrium on which coordinate.[25] The notion that agents might coordinate on such variables is not so far-fetched. Roos (2008), for instance, provides survey evidence showing that students overweight realizations of nonfundamental factors relative to more fundamental factors in assessing the impacts of those factors on short-run macroeconomic performance in Germany. However, without the controlled conditions of the laboratory, it can be difficult to say what factors are truly fundamental, which are less so, and which are purely extrinsic and nonfundamental. Three experimental studies of sunspot variables as coordination devices have been conducted: Marimon and others (1993), Duffy and Fisher (2005), and Fehr, Heinemann, Llorente-Saguer (2013); we describe each in turn.

Marimon and Sunder (1993) implemented a two-period overlapping generations environment where, if agents have perfect foresight, there are multiple equilibria: an interior steady state and a two-period cyclic equilibrium. Subjects in the role of young agents formed price expectations that determined current prices, given the nonlinear model, $p_t = \phi(p^e_{t+1})$. Thus given price expectations, subjects' optimal consumption and savings in the form of real money balances were determined (as in Marimon and Sunder 1993, 1994). Marimon and Sunder hoped that subjects would use realizations of a sunspot variable to coordinate their expectations on the cyclic equilibrium. Their sunspot variable consisted of a blinking cube on subjects' computer screens. The color of this cube alternated every period between red and yellow. Marimon and Sunder found that subjects essentially ignored the sunspot variable realizations and simply coordinated on the steady states. They later tried to add a correlation between the sunspot variable and a real endowment shock (alternating the size of the young generation between three and four subjects, i.e., three-four-three-four), but this also did not lead to coordination on the sunspot variable when the endowment shock was shut off.

Duffy and Fisher (2005) consider a simpler, partial equilibrium framework that abstracts from a number of conceptual difficulties (e.g., implementing an infinite horizon). In this simple and static environment, there are two equilibria that differ only in terms of the equilibrium price level; the equilibrium quantity is the same in both. The experimental design involves five buyers and five sellers, each with two units to buy or sell. Buyers seek to maximize consumer surplus (valuation − price), while sellers seek to maximize producer's surplus (price − cost). Further, each buyer (seller) had two possible valuations (costs) for each of his or her two units. If the state was high, each buyer's (seller's) profits were calculated using his or her two high valuations (costs). If the state was low, each buyer's (seller's) profits were calculated using his or her two low valuations (costs). The two sets of valuations/costs used in the experiment are shown in Figure 1.5. Buyers are B1–B5 and sellers are S1–S5. Market clearing prices with high demand and supply are in the interval [190, 210]. Market clearing prices with low demand and supply are in the interval [90, 110]. The equilibrium quantity is always six units bought and sold.

Two market-clearing mechanisms were considered—the standard double auction, where bids and asks can be observed in real time, and a sealed-bid variant known as a call market, where bids and asks are submitted simultaneously, bids are sorted from highest to lowest and asks, from lowest to highest, and a single market clearing price is determined by the intersection of demand and supply (if there is one). All buyers with bids above the market price get to buy their units, provided there are enough units for sale. All sellers with asks below the market price get to sell their units provided their is enough demand. The state of the world was determined by the median traded price in the double auction or by the market-clearing price in the call market. If either price was greater than or equal to 150, then the high state was declared, and subjects use high valuations or costs in determining their surplus (payoff). Otherwise the low state was declared, and low valuations and costs were used in the determination of payoffs. Thus the situation is akin to one in which there are multiple equilibria, each supported by different beliefs about the likely state of the world.

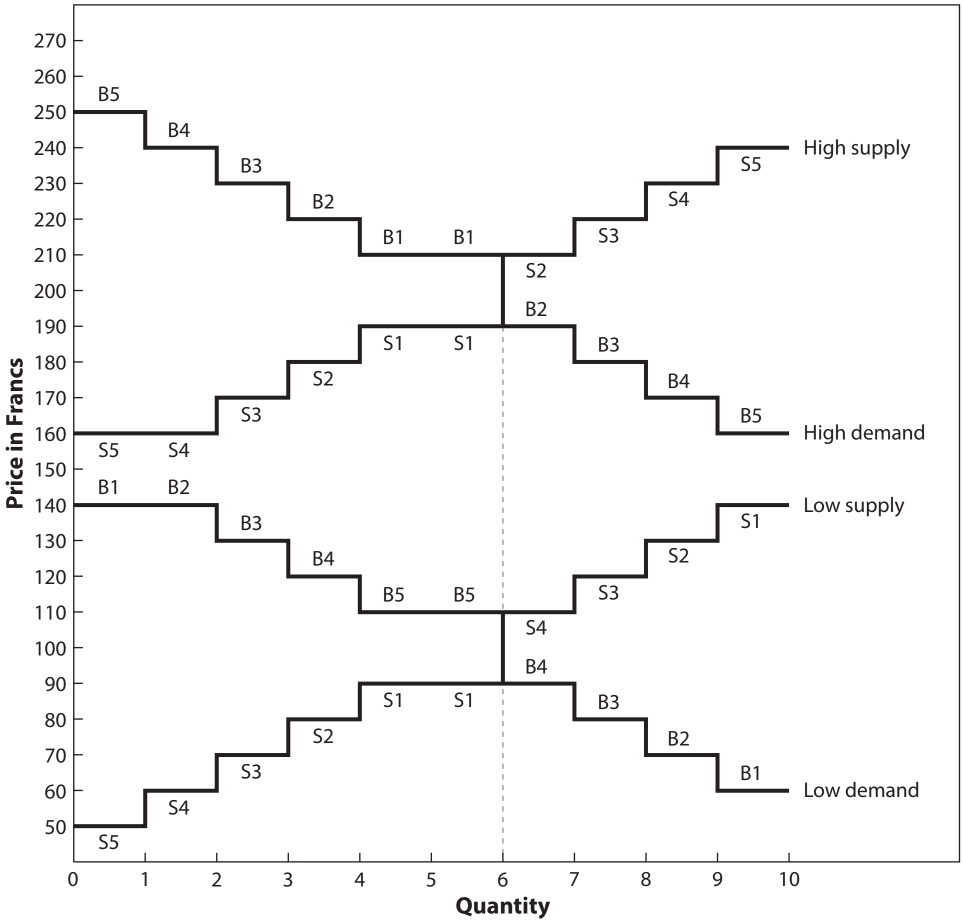

Figure 1.5: Induced high and low demand and supply in Duffy and Fisher (2005).

Duffy and Fisher's sunspot variable was one of two possible announcements made prior to each of ten four-minute trading periods. The announcement chosen was determined by publicly flipping a coin. In one treatment, if the coin flip was heads, the public announcement was, "the forecast is high", while if the coin flip was tails, the public announcement was, "the forecast is low", and this scheme was public knowledge. Duffy and Fisher report that in sessions using a call-market-clearing mechanism, subjects perfectly coordinated on the high-price equilibrium when the forecast was high and on the low-price equilibrium when the forecast was low—that is, the sunspot variable was shown to matter for economic volatility. On the other hand, under the double-auction-market-clearing mechanism, the sunspot announcements only sometimes served to coordinate subjects on the high or low equilibrium. Duffy and Fisher argue that the reason for this difference lies in the real-time information that was available in the double auction; subjects could see bids and asks as they occurred and could use this information to attempt to engineer an equilibrium outcome for prices (high or low) that was more favorable to them.[26] Thus the coordinating mechanism provided by the sunspot could be undone by the real-time information on bids, asks, and trade prices. The same was not possible in the call market, where bids and asks had to be submitted simultaneously; hence the sunspot variable played an important coordinating role in the environment.

Duffy and Fisher further show that the semantics of the sunspot variable matter: replacing "the forecast is high" or "low" with "the forecast is sunshine" or "rain" eliminated the sunspot variable as a coordinating mechanism in the call market.

Fehr, Heinemaan, and Llorente-Saguer (2013) study the emergence of sunspot equilibria in an even simpler setting, a two-player coordination game, where the two players i (j) must simultaneously choose numbers a_i (a_j), from the interval [0, 100], and each earns a payoff that is a quadratic function of the squared deviation, $(a_i - a_j)^2$. The focus of this study is on the nature and number of the extrinsic signals, whether they must be public or could be privately observed and whether there is one signal or two. In most treatments a common extrinsic signal, Z, is known to be a random drawn from the binary distribution {0, 100} at the start of each of 80 periods. In some treatments, the value of Z is publicly observable to both players, while in other treatments subjects receive a private noisy signal of the value of Z with a given precision, or a public and private signal, or two public signals, all from the same binary distribution. In a control treatment, subjects receive no signal and quickly coordinate on the risk-dominant choice of 50 (the midpoint of the action space). When there is a single public signal, subjects play according to a sunspot equilibrium, choosing numbers corresponding to the realized public signal 0 or 100. They have no difficulty continuing to play according to a sunspot equilibrium with two public signals; when the signals differ, they choose the average of the two signals, 50, and thus coordinate on play of a "three-cycle." The sunspot equilibrium breaks down when subjects receive a public and a private signal, as subjects are unable to ignore their private signal and, consequently, their play converges to the risk dominant strategy of always choosing 50. Most interestingly, they report that if subjects receive only private signals of Z (no public signal) and these private signals are sufficiently precise as to the true value of Z so that the private signals are highly correlated with one another, then subjects continued to choose numbers according to the private signal they received even though such actions are not consistent with any pure sunspot equilibrium. This is an interesting empirical finding, suggesting an avenue by which the notion of a sunspot equilibrium might be more general than theory currently admits.

Further research on this topic might seek to understand how the mapping from sunspot variable realizations to the action space matters in getting subjects to coordinate on sunspot equilibria; for instance, does the dimensionality of the signal space need to be small relative to the action space, and if so, how small? It would also be of interest to consider sunspot equilibria that are not simply randomizations over two certainty equilibria.

3.4 Resolving Coordination Problems: The Global Game Approach

Another view of multiple equilibria in macroeconomic modeling is that the equilibrium beliefs in support of these equilibria may not be as indeterminate as theory supposes. As Morris and Shin (2001) argue, these indeterminacies arise from assuming that economic fundamentals are *common knowledge* and that individuals are certain of the behavior of others in equilibrium. Relaxing these assumptions—for example, by introducing some uncertainty about fundamentals—can remove the multiplicity, à la the Carlsson and van Damme's (1993) global game approach for 2×2 games.[27] The resulting game is one in which individuals adopt a unique threshold strategy—when fundamentals are

(continued...)

NAME INDEX

SUBJECT INDEX