CONTENTS

1

Macroeconomics: A Survey of Laboratory Research

John Duffy

1 INTRODUCTION: LABORATORY MACROECONOMICS

Macroeconomic theories have traditionally been tested using nonexperimental field data, most often national income account data on GDP and its components. This practice follows from the widely held belief that macroeconomics is a purely observational science: history comes around just once and there are no "do-overs." Controlled manipulation of the macroeconomy to gain insight regarding the effects of alternative institutions or policies is viewed by many as impossible, not to mention unethical, and so, apart from the occasional *natural* experiment, most macroeconomists would argue that macroeconomic questions cannot be addressed using experimental methods.[1]

Yet, as this survey documents, over the past twenty-five years, a wide variety of macroeconomic models and theories *have* been examined using controlled laboratory experiments with paid human subjects, and this literature is growing. The use of laboratory methods to address macroeconomic questions has come about in large part due to changes in macroeconomic modeling, though it has also been helped along by changes in the technology for doing laboratory experimentation, especially the use of large computer laboratories. The change in macroeconomic modeling is, of course, the now widespread use of explicit microfounded models of constrained, intertemporal choice in competitive general equilibrium, game-theoretic or search-theoretic frameworks. The focus of these models is often on how institutional changes or policies affect the choices of decision makers such as household and firms, in addition to the more traditional concern with responses in the aggregate time series data (e.g., GDP) or to the steady states of the model. While macroeconomic models are often expressed at an aggregate level—for instance, there is a "representative" consumer or firm or a market for the "capital good"—an implicit, working assumption of many macroeconomists is that aggregate sectoral behavior is not different from that of the individual actors or components that comprise each sector.[2] Otherwise, macroeconomists would be obliged to be explicit about the mechanisms by which individual choices or sectors aggregate up to the macroeconomic representations they work with, and macroeconomists have been largely silent on this issue. Experimentalists testing nonstrategic macroeconomic models have sometimes taken this representativeness assumption at face value and conducted individual decision-making experiments with a macroeconomic flavor. But, as we shall see, experimentalists have also considered whether small groups of subjects interacting with one another via markets or by observing or communicating with one another might outperform individuals in tasks that macroeconomic models assign to representative agents.

While there is now a large body of macroeconomic experimental research as reviewed in this survey, experimental methods are not yet a mainstream research tool used by the typical macroeconomist, as they are in nearly every other field of economics. This state of affairs likely arises from the training that macroeconomists receive, which does not typically include exposure to laboratory methods and is instead heavily focused on the construction of dynamic stochastic general equilibrium models that may not be well suited to experimental testing. As Sargent (2008, p 27) observes,

> I suspect that the main reason for fewer experiments in macro than in micro is that the choices confronting artificial agents within even one of the simpler recursive competitive equilibria used in macroeconomics are very complicated relative to the settings with which experimentalists usually confront subjects.

This complexity issue can be overcome, but, as we shall see, it requires experimental designs that simplify macroeconomic environments to their bare essence or involve operational issues such as the specification of the mechanism used to determine equilibrium prices. Despite the complexity issue, I will argue in this survey that experimental methods can and should serve as a complement to the modeling and empirical methods currently used by macroeconomists as laboratory methods can shed light on important questions regarding the empirical relevance of microeconomic foundations, questions of causal inference, equilibrium selection and the role of institutions.[3]

Indeed, to date the main insights from macroeconomic experiments include (1) an assessment of the microassumptions underlying macroeconomic models, (2) a better understanding of the dynamics of forward-looking expectations, which play a critical role in macroeconomic models, (3) a means of resolving equilibrium selection (coordination) problems in environments with multiple equilibria, (4) validation of macroeconomic model predictions for which the relevant field data are not available, and (5) the impact of various macroeconomic institutions and policy interventions on individual behavior. In addition, laboratory tests of macroeconomic theories have generated new or strengthened existing experimental methodologies, including implementation of the representative-agent assumption, overlapping generations, and search-theoretic models, methods for assisting with the roles of forecasting and optimizing, implementation of discounting and infinite horizons, methods for assessing equilibration, and the role played by various market clearing mechanisms in characterizing Walrasian competitive equilibrium (for which the precise mechanism of exchange is left unmodeled).

The origins of *macroeconomic experiments* are unclear. Some might point to A. W. Phillips' (1950) experiments using a colored liquid-filled tubular flow model of the macroeconomy, though this did not involve human subjects! Others might cite Vernon Smith's (1962) double-auction experiment demonstrating the importance of centralized information to equilibration to competitive equilibrium as the first macroeconomic experiment. Yet another candidate might be John Carlson's (1967) early experiment examining price expectations in stable and unstable versions of the cobweb model. However, I will place the origins more recently with Lucas's

1986 invitation to macroeconomists to conduct laboratory experiments to resolve macrocoordination problems that were unresolved by theory. Lucas's invitation was followed up on by Aliprantis and Plot (1992), Lim, Prescott, and Sunder (1994), and Marimon and Sunder (1993, 1994, 1995), and, perhaps as the result of their interesting and influential work, over the past two decades, there has been a great blossoming of research testing macroeconomic theories in the laboratory. This literature is now so large that I cannot hope to cover every paper in a single chapter, but I do hope to give the reader a good road map as to the kinds of macroeconomic topics that have been studied experimentally as well as to suggest some further extensions.

How shall we define a macroeconomic experiment? One obvious dimension might be to consider the number of subjects in the study. Many might argue that a macroeconomic experiment should involve a *large* number of subjects; perhaps the skepticism of some toward macroeconomic experiments has to do with the necessarily small numbers of subjects (and small scale of operations) that are possible in laboratory studies.[4] The main problem with small numbers of subjects is that strategic considerations may play a role that is not imagined (or possible) in the macroeconomic model that is being tested, which may instead focus on perfectly competitive Walrasian equilibrium outcomes. However, research has shown that attainment of competitive equilibrium outcomes might not require large numbers of subjects. For example, the evidence from numerous double-auction experiments beginning with Smith (1962) and continuing to the present reveals that equilibration to competitive equilibrium can occur reliably with as few as three to five buyers or sellers on each side of the market. Duffy and others (2011) study bidding behavior in a Shapley-Shubik market game and show that with small numbers of subjects (e.g., groups of size two), Nash equilibrium outcomes are indeed far away from the competitive equilibrium outcome of the associated pure exchange economy. However, they also show that as the number of subjects increases, the Nash equilibrium subjects coordinate upon becomes approximately Walrasian; economies with just ten subjects yield market-based allocations that are indistinguishable from the competitive equilibrium of the associated pure exchange economy. Thus, while more subjects are generally better than fewer subjects for obtaining competitive equilibrium outcomes, it seems possible to establish competitive market conditions with the small numbers of subjects available in the laboratory.[5]

A more sensible approach is to define a macroeconomic experiment as one that tests the predictions of a macroeconomic model or its assumptions or is framed in the language of macroeconomics, involving, for example, intertemporal consumption and savings decisions, inflation and unemployment, economic growth, bank runs, monetary exchange, monetary or fiscal policy, or any other macroeconomic phenomena. Unlike microeconomic models and games, which often strive for generality, macroeconomic models are typically built with a specific macroeconomic story in mind that is not as easily generalized to other nonmacroeconomic settings. For this reason, our definition of a macroeconomic experiment may be too restrictive. There are many microeconomic experiments—coordination games for instance—that can be given both a macroeconomic interpretation or a more microeconomic interpretation, for example; as models of firm or team behavior. In discussing those studies as macroeconomic experiments, I will attempt to emphasize the macroeconomic interpretation.

The coverage of this chapter can be viewed as an update on some topics covered in several chapters of the first volume of the *Handbook of Experimental Economics*, including discussions of intertemporal decision making by Camerer (1995), coordination problems by Ochs (1995), and asset prices by Sunder (1995), though the coverage here will not be restricted to these topics alone. Most of the literature surveyed here was published since 1995, the date of the first *Handbook* volume. In addition, this chapter builds on, complements, and extends earlier surveys of the macroeconomic experimental literature by myself, Duffy (1998, 2008), and by Ricciuti (2008).

2 DYNAMIC, INTERTEMPORAL OPTIMIZATION

Perhaps the most widely used model in modern macroeconomic theory is the one-sector, infinite-horizon optimal-growth model pioneered by Ramsey (1928) and further developed by Cass (1965) and Koopmans (1965). This model posits that individuals solve a dynamic, intertemporal optimization problem in deriving their consumption and savings plan over an infinite horizon. Both deterministic and stochastic versions of this model are workhorses of modern real business cycle theory and growth theory.

In the urge to provide microfoundations for macroeconomic behavior, modern macroeconomists assert that the behavior of consumers or firms can be reduced to that of a representative, fully rational individual actor; there is no room for any "fallacies of composition" in this framework. It is, therefore, of interest to assess the extent to which macroeconomic phenomena can be said to reflect the choices of individuals facing dynamic stochastic intertemporal optimization problems. Macroeconomists have generally ignored the plausibility of this choice-theoretic assumption, preferring instead to examine the degree to which the time-series data on GDP and its components move in accordance with the conditions that have been optimally derived from the fully rational representative-agent model and especially whether these data react predictably to shocks or policy interventions.

2.1 *Optimal Consumption/Savings Decisions*

Whether individuals can in fact solve a dynamic stochastic intertemporal optimization problem of the type used in the one-sector optimal growth framework has been the subject of a number of laboratory studies, including Hey and Dardanoni (1988), Carbone and Hey (2004), Noussair and Matheny (2000), Lei and Noussair (2002), Ballinger, Palumbo, and Wilcox (2003), Carbone (2006), Brown, Chua, and Camerer (2009), Ballinger and others (2011), Crockett and Duffy (2013), Carbone and Duffy (2014), and Meissner (2016), among others. These studies take the representative agent assumption of modern macroeconomics seriously and ask whether subjects can solve a discrete-time optimization problem of the form

$$\max_{\{c_t\}} E_t \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$c_t + x_t \leq \omega_t$$

where c_t is time t consumption, $u(\cdot)$ is a concave utility function, β is the period discount factor, x_t represents time t savings (if positive) or borrowings (if negative), and ω_t is the household's time t wealth.

Hey and Dardanoni (1988) assume a pure exchange economy, where wealth evolves according to $\omega_t = R(\omega_{t-1} - c_{t-1}) + y_t$, with $\omega_0 > 0$ given. Here, R denotes the (constant) gross return on savings and y_t is the stochastic time t endowment of the single good; the mean and variance of the stochastic income process is made known to subjects. By contrast, Noussair and associates assume a nonstochastic production economy, where $\omega_t = f(k_t) + (1-\delta)k_t$, with $f(\cdot)$ representing the known, concave production function, k_t denoting capital per capita, and δ denoting the depreciation rate. In this framework, it is public knowledge that an individual's savings, x_t, are invested in capital and become the next period's capital stock, that is, $x_t = k_{t+1}$. The dynamic law of motion for the production economy is expressed in terms of capital rather than wealth: $k_{t+1} = f(k_t) + (1-\delta)k_t - c_t$, with $k_0 > 0$ given. The gross return on savings is endogenously determined by $R = f'(k_t) + (1-\delta)$.

Solving the maximization problem given before, the first-order conditions imply that the optimal consumption program must satisfy the Euler equation

$$u'(c_t) = \beta R E_t u'(c_{t+1})$$

where the expectation operator is with respect to the (known) stochastic process for income (or wealth). Notice that the Euler equation predicts a monotonic increasing, decreasing, or constant consumption sequence, depending on whether βR is less than, greater than, or equal to 1. Solving for a consumption or savings *function* involves application of dynamic programming techniques that break the optimization problem up into a sequence of two-period problems; the Euler equation characterizes the dynamics of marginal utility in any two periods. For most specifications of preferences, analytic closed-form solutions for the optimal consumption or savings function are not possible, though via concavity assumptions, the optimal consumption/savings program can be shown to be unique.

In testing this framework, Hey and Dardanoni (1988) addressed several implementation issues. First, they chose to rule out borrowing (negative saving) in order to prevent subjects from ending the session in debt. Second, they attempted to implement discounting and the stationarity associated with an infinite horizon by having a constant probability that the experimental session would continue with another period.[6] Finally, rather than inducing a utility function, they supposed that all subjects had constant absolute risk-aversion preferences, and they estimated each individual subject's coefficient of absolute risk aversion using data they gathered from hypothetical and paid choice questions presented to the subjects. Given this estimated utility function, they then numerically computed optimal consumption for each subject and compared it with their actual consumption choice. To challenge the theory, they considered different values for R and β as well as for the parameters governing the stochastic income process, y.

They report mixed results. First, consumption is significantly different from optimal behavior. In particular, there appears to be great time dependence in consumption behavior; that is, consumption appears dependent on past income realizations, which is at odds with the time-independent nature of the optimal consumption program. Second, they find support for the comparative statics implications of the theory. That is, changes in the discount factor, β, or in the return on savings, R, have the same effect on consumption as under optimal consumption behavior. So they find mixed support for dynamic intertemporal optimization.

Carbone and Hey (2004) and Carbone (2006) simplify the design of Hey and Dardanoni. First, they eliminate discounting and consider a finite-horizon, twenty-five period model. They argue, based on the work of Hey and Dardanoni, that subjects "misunderstand the stationarity property" of having a constant probabilistic stopping rule. Second, they greatly simplify the stochastic income process, allowing there to be

TABLE 1.1:
Average change in consumption in response to parameter changes and conditional on employment status. Source: Carbone and Hey (2004, Table 5).

Change (Δ) in Treatment Variable (from low value to high value)	*Unemployed*		*Employed*	
	Optimal	*Actual*	*Optimal*	*Actual*
Δp (Pr. remaining employed)	5.03	23.64	14.57	39.89
Δq (Pr. becoming employed)	14.73	−1.08	5.68	0.15
Δ Ratio high-low income	0.25	0.24	0.43	0.76

just two values for income—one high, which they refer to as a state where the consumer is "employed," and the other low, in which state the consumer is "unemployed." They use a two-state Markov process to model the state transition process: conditional on being employed (unemployed), the probability of remaining (becoming) employed was $p(q)$, and these probabilities were made known to subjects. Third, rather than infer preferences they induce a constant absolute risk-aversion utility function. Their treatment variables were p, q, R and the ratio of employed to unemployed income; they considered two values of each, one high and one low, and examined how consumption changed in response to changes in these treatment variables relative to the changes predicted by the optimal consumption function (again numerically computed). Table 1.1, shows a few of their comparative statics findings.

An increase in the probability of remaining employed caused subjects to overreact in their choice of additional consumption relative to the optimal change regardless of their employment status (unemployed or employed), whereas an increase in the probability of becoming employed—a decrease in the probability of remaining unemployed—led to an underreaction in the amount of additional consumption chosen relative to the optimal prediction. On the other hand, the effect of a change in the ratio of high-to-low income on the change in consumption was quite close to optimal. Carbone and Hey emphasize also that there was tremendous heterogeneity in subjects' abilities to confront the life-cycle consumption savings problem, with most subjects appearing to discount old-age consumption too heavily (when they should not discount at all) or optimizing over a shorter planning horizon than the twenty-five periods of the experiment.[7] Carbone and Hey conclude that "subjects do not seem to be able to smooth their consumption stream sufficiently—with current consumption too closely tracking current income." Interestingly, the *excess sensitivity* of consumption to current income (in excess of that warranted by a revision in expectations of future income) is a well-documented empirical phenomenon in studies of consumption behavior using aggregate field data (see, e.g., Flavin 1981; Hayashi 1982; Zeldes 1989). This corroboration of evidence from the field should give us further confidence in the empirical relevance of the laboratory analyses of intertemporal consumption-savings decisions. Two explanations for the excess sensitivity of consumption to income that have appeared in the literature are (1) binding liquidity constraints and (2) the presence of a precautionary savings motive (which is more likely in a finite-horizon model). Future experimental research might explore the relative impacts of these two factors on consumption decisions.

Meissner (2016) modifies the finite-horizon, life-cycle planning environment of Carbone and Hey (2004) to allow subjects to borrow and not just to save. In particular, Meissner studies two regimes, one in which an individual's stochastic income process has an upward-sloping trend and a second regime where this income process has a downward-sloping trend. Optimal behavior in the first regime involves borrowing in the early periods of life so as to better smooth consumption, while optimal behavior in the second regime involves saving in the early periods of life to better smooth consumption. Meissner parameterized the environment so that the optimal consumption path was the same in both income treatments, and subjects were given three opportunities or "lifetimes" to make consumption/savings/borrowing decisions in each of the two income treatments, that is, he uses a within-subjects design. A main finding is that in the decreasing-income regime, subjects have no trouble learning to save in the early periods of their life and can approximately smooth consumption over their lifetime. By contrast, in the increasing-income regime, most subjects seem averse to borrowing any amount, so that consumption deviates much further from the optimal path; consumption decisions in this treatment more closely track the upward-trend path of income and there is not much difference with replication (i.e., there is little learning). Meissner attributes the latter finding to "debt aversion" on the part of his university student subjects. It would be of interest to explore whether such debt aversion continues in more-general subject populations involving individuals who may have some homegrown experience with acquiring debt.

Noussair and Matheny (2000) further modify the framework of Hey and associates by adding a concave production technology, $f(k_t) = Ak_t^{\alpha}$, $\alpha < 1$, which serves to endogenize the return on savings in conformity with modern growth theory. They induce both the production function and a logarithmic utility function by giving subjects schedules of payoff values for various levels of k and c, and they implement an infinite horizon by having a constant probability that a sequence of rounds continues. Subjects made savings decisions (chose $x_t = k_{t+1}$) with the residual from their budget constraint representing their consumption. Noussair and Matheny varied two model parameters, the initial capital stock k_0 and the production function parameter α. Variation in the first parameter changes the direction by which paths for consumption and capital converge to steady-state-values (from above or below) while variations in the second parameter affect the predicted *speed* of convergence; the lower is α, the greater is the speed of convergence of the capital stock and consumption to the steady state of the model. Among the main findings, Noussair and Matheny report that sequences for the capital stock are monotonically decreasing regardless of parameter conditions, and theoretical predictions with regard to speed of convergence do not find much support. Consumption is, of course linked to investment decisions and is highly variable. They report that subjects occasionally resorted to consumption *binges*, allocating nearly nothing to the next period's capital stock in contrast to the prediction of consumption smoothing. However, this behavior seemed to lessen with experience. A virtue of the Noussair-Matheny study is that it was conducted with both US and Japanese subjects, with similar findings for both countries.

One explanation for the observed departure of behavior from the dynamically optimal path is that the representative-agent assumption, while consistent with the reductionist view of modern macroeconomics, assumes too much individual rationality to be useful in practice.[8] Information on market variables (e.g., *prices*) as determined by *many different interacting agents*, may be a necessary aid to solving such complicated optimization decisions. An alternative explanation may be that the standard model of intertemporal consumption smoothing abstracts away from the importance of *social norms* of behavior with regard to consumption decisions. Akerlof (2007), for instance, suggests that people's consumption decisions may simply reflect their "station in life." College students (the subjects in most of these experiments) looking to their peers, choose to *live like college students* with expenditures closely tracking income. Both of these alternative explanations have been considered to some extent in further laboratory studies.

Crockett and Duffy (2013) explore whether groups of subjects can learn to intertemporally smooth their consumption in the context of an infinite-horizon, consumption-based asset-pricing model, specifically, the Lucas tree model (Lucas 1978). In the environment they study, the only means of saving intertemporally is to buy or sell shares of a long-lived asset (a Lucas tree), which yields a known and constant divided (amount of fruit) each period. Subjects are of two types, according to the endowment of income they receive in alternating periods; odd types receive high income in odd-numbered periods and low income in even-numbered periods, while even types receive high income in even-numbered periods and low income in odd-numbered periods. In one of Crockett and Duffy's treatments, subjects' induced utility function over consumption is concave so that subjects have an incentive to intertemporally smooth their consumption by buying the asset in their high-income periods and selling it in their low-income periods (the heterogeneity of subject types allows for such trades to occur). Asset prices are determined via a double-auction mechanism, and these prices can be observed by all subject participants. Crockett and Duffy report that with these asset price signals, most subjects have little difficulty learning to intertemporally smooth their consumption across high- and low-income periods. Future experimental research on consumption smoothing through the purchase and sale of long-lived assets might investigate a more realistic, stochastic, life-cycle income process.

Ballinger, Palumbo, and Wilcox (2003) explore the role of social learning in a modified version of the noisy pure exchange economy studied by Hey and Dardanoni (1988). In particular, they eliminate discounting (presumably to get rid of time dependence), focusing on a finite sixty-period horizon. Subjects are matched into three-person "families" and make decisions in a fixed sequence. The generation 1 (G1) subject makes consumption decisions alone for twenty periods; in the next twenty periods (21–40), his or her behavior is observed by the generation 2 (G2) subject, and in one treatment, the two are free to communicate with one another. In the next twenty periods (periods 41–60 for G1, periods 1–20 for G2), both generations make consumption/savings decisions. The G1 subject then exits the experiment. The same procedure is then repeated with the generation 3 (G3) subject watching the G2 subject for the next twenty rounds, and so on. Unlike Hey and Dardanoni, Ballinger and others induce a constant relative risk-aversion utility function on subjects using a Roth and Malouf (1979) binary lottery procedure. This allows them to compute the path of optimal consumption/savings behavior. These preferences give rise to a precautionary savings motive, wherein liquid wealth (saving) follows a hump-shaped pattern over the sixty-period lifecycle.

Ballinger, Palumbo, and Wilcox's (2003) main treatment variable concerns the variance of the stochastic income process (high or low), which affects the peak of the precautionary savings hump; in the high case they also explore the role of allowing communication/mentoring or not (while maintaining observability of actions by overlapping cohorts at all times). Among their findings, they report that subjects tend to consume more than the optimal level in the early periods of their lives, leading to less savings and below-optimal consumption in the later periods of life. However, savings are greater in the high- as compared with the low-variance case, which is consistent with the comparative statics prediction of the rational intertemporal choice framework. They also find evidence for time dependence in that consumption behavior is excessively sensitive to near lagged changes in income. Most interestingly, they report that consumption behavior of generation 3 is significantly closer to the optimal consumption program than in the consumption behavior of generation 1, suggesting that social learning by observation plays an important role and may be a more reasonable characterization of the representative agent.

Ballinger and others (2011) study a similar life-cycle consumption/savings problem but focus on whether cognitive and/or personality measures might account for the observed heterogeneity in a subject's savings behavior, in particular, the subject's use of shorter-than-optimal planning horizons. Using a careful multivariate regression analysis that accounts for potentially confounding demographic variables, they report that cognitive measures and not personality measures are good predictors of heterogeneity in savings behavior. In particular, they report that variations in subjects' cognitive abilities, as assessed, using visually oriented "pattern-completion" tests and "working memory" tests that assess a subject's ability to control both attention and thought, can explain variations in a subject's life-cycle savings behavior and that the median subject is thinking just three periods ahead.

Lei and Noussair (2002) study the intertemporal consumption savings problem in the context of the one-sector optimal growth model with productive capital. They contrast the "social planner" case, where a single subject is charged with maximizing the representative consumer firms' present discounted sum of utility from consumption over an indefinite horizon (as in Noussair and Matheny (2000)), with a decentralized market approach, wherein the same problem is solved by five subjects looking at price information. In this market treatment, the production and utility functions faced by the social planner are disaggregated into five individual functions assigned to the five subjects that aggregate up to the same functions faced by the social planner. For example, some subjects had production functions with marginal products for capital that were higher than for the economy-wide production function, while others had marginal products for capital that are lower. At the beginning of a period, production took place, based on previous period's capital, using either the individual production functions in the market treatment or the economy-wide production function in the social-planner treatment. Next, in the market treatment, a double-auction market for output (or potential future capital) opened up. Agents with low marginal products of capital could trade some of their output to agents with high marginal products for capital in exchange for experimental currency units (subjects were given an endowment of such units each period, which they had to repay). The import of this design was that the market effectively communicated to the five subjects the market price of a unit of output (or future capital). As future capital could be substituted one for one with future consumption, the market price of capital revealed to subjects the marginal utility of consumption. After the market for output closed, subjects in the market treatment could individually allocate their adjusted output levels between future capital k_{t+1} or savings and experimental currency units or consumption c_t. By contrast, in the social-planner treatment, there was no market for output; the representative individual proceeded directly to the step of deciding how to allocate output between future capital (savings) and current consumption. At the end of the period, subjects' consumption amounts were converted into payoffs using the economy-wide or individual concave utility functions, and loans of experimental currency units in the market treatment were repaid.

The difference in consumption behavior between the market and representative-agent–social planner treatments is illustrated in Figure 1.1, which shows results from a representative session of one of Lei and Noussair's treatments. In the market treatment,

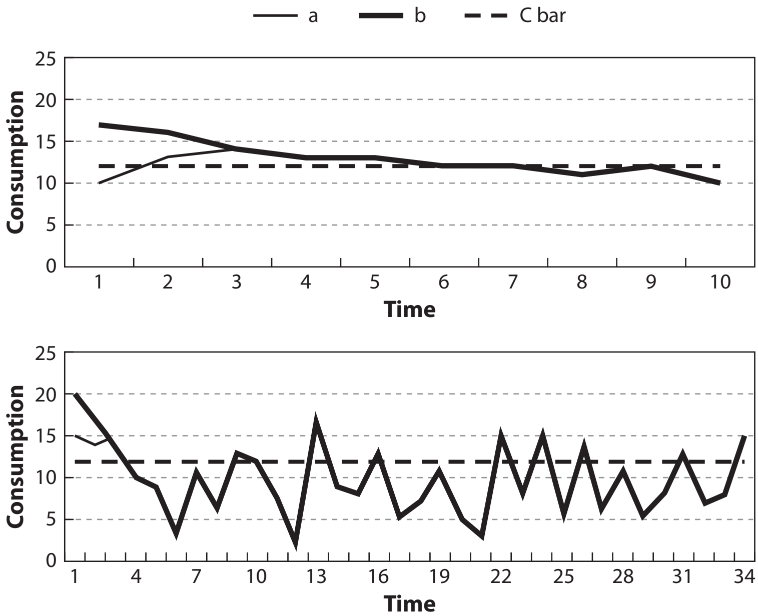

Figure 1.1: Consumption choices over two indefinite horizons (a, b) compared with optimal steady-state consumption (C bar). Market treatment (*top*) versus social planner treatment (*bottom*). Source: Lei and Noussair (2002).

there was a strong tendency for consumption (as well as capital and the price of output) to converge to their unique steady-state values, while in the social planner treatment, consumption was typically below the steady-state level and much more volatile.

In further analysis, Lei and Noussair (2002) make use of a linear, panel data-regression model to assess the extent to which consumption and savings (or any other time-series variable for that matter) can be said to be converging over time toward predicted (optimal) levels.[9] In this regression model, $y_{j,t}$ denotes the average (or economy-wide level) of the variable of interest by cohort/session j in period $t = 1, 2, \ldots,$ and D_j is a dummy variable for each of the $j = 1, 2, \ldots, J$ cohorts. The regression model is written as

$$y_{j,t} = \alpha_1 \frac{D_1}{t} + \alpha_2 \frac{D_2}{t} \cdots + \alpha_J \frac{D_J}{t} + \beta \frac{t-1}{t} + \epsilon_{j,t} \quad (1)$$

where $\epsilon_{j,t}$ is a mean zero, random error term. The α_j coefficients capture the initial starting values for each cohort, while the β coefficient captures the asymptotic value of the variable y to which all J cohorts of subjects are converging; notice that the α coefficients have a full weight of 1 in the initial period 1 and then have exponentially declining weights, while the single β coefficient has an initial weight of zero that increases asymptotically to 1. For the dependent variable in (1), Lei and Noussair (2002) use: (1) the consumption and capital stocks (savings) of cohort j, $c_{j,t}$, and $k_{j,t+1}$, (2) the absolute deviation of consumption from its optimal steady-state value, $|c_{j,t} - c^*|$, and (3) the ratio of the realized utility of consumption to the optimum, $u(c_{j,t})/u(c^*)$. For the first type of dependent variable, the estimate $\hat{\beta}$ reveals the values to which the dependent variable, $c_{j,t}$ and $k_{j,t}$ are converging across cohorts; strong convergence is said to obtain if $\hat{\beta}$ is not significantly different from the optimal steady-state levels, c^* and k^*. For the second and third types of dependent variable, one looks for whether $\hat{\beta}$ is significantly different from zero or one, respectively. Lei and Noussair also consider a weaker form of convergence that examines whether $\hat{\beta}$ is closer (in absolute value) to the optimal, predicted level than a majority of the $\hat{\alpha}_j$ estimates. Using all four dependent variables, they report evidence of both weak and strong convergence in the market treatment, but only evidence of weak (and not strong) convergence in the social planner treatment.[10]

Tests of convergence based on the regression model (1) can be found in several experimental macroeconomic papers reviewed later in this chapter. This methodology for assessing convergence of experimental time series is one of several methodologies that might be considered "native" to experimental macroeconomics. Therefore, allow me a brief digression on the merits of this approach. First, the notion that strong convergence obtains if $\hat{\beta}$ is not significantly different from the predicted level, y^*, while weak convergence obtains if $|\hat{\beta} - y^*| < |\hat{\alpha}_j - y^*|$ for a majority of js is somewhat problematic, as strong convergence need not imply weak convergence, as when the $\hat{\alpha}_j$ estimates are insignificantly different from $\hat{\beta}$. Second, if *convergence* is truly the focus, an alternative approach would be to use an explicitly *dynamic* adjustment model for each cohort j of the form

$$y_{j,t} = \lambda_j y_{j,t-1} + \mu_j + \epsilon_{j,t} \tag{2}$$

Using (2), *weak* convergence would obtain if the estimates, $\hat{\lambda}_j$, were significantly less than 1, while *strong* convergence would obtain if the estimate of the long-run expected value for y_j, $\hat{\mu}_j/(1-\hat{\lambda}_j)$, was not significantly different from the steady-state prediction y^*; in this model, strong convergence implies weak convergence, not the reverse.[11] Finally, analysis of joint convergence across the J cohorts to the predicted level y^* could be studied through tests of the hypothesis:

$$I_J \begin{pmatrix} \hat{\mu}_1 \\ \vdots \\ \hat{\mu}_J \end{pmatrix} + \begin{pmatrix} \hat{\lambda}_1 \\ \vdots \\ \hat{\lambda}_J \end{pmatrix} y^* = \begin{pmatrix} y^* \\ \vdots \\ y^* \end{pmatrix}$$

where I_J is a J-dimensional identity matrix.

Returning to the subject of dynamic, intertemporal life-cycle consumption/savings decisions, recent work has explored subject behavior in the case where there are *two* (as opposed to just one) state variables: an individual's wealth (or "cash on hand") ω_t and some induced "habit" level for consumption h_t (following the macroeconomic literature on habit formation), so that the period objective function is of the form $u(c_t, h_t)$. Brown, Chua, and Camerer (2009) study the case of *internal* habit formation, where each individual subject i has his or her own, personal habit level of consumption that evolves according to $h_t^i = \alpha h_{t-1}^i + c_t^i$ ($\alpha < 1$) and has a period utility function that is increasing in the ratio of c_t^i/h_t^i. Carbone and Duffy (2014) study the case of *external* habit formation, where h_t is the lagged average consumption of a group of N identically endowed subjects (i.e., $h_t = N^{-1}\sum_{i=1}^N c_{t-1}^i$) and u is an increasing function of the difference, $c_t - \alpha h_t$ ($\alpha < 1$). Both studies also explore social learning in this more complex environment, with Brown and others exploring intergenerational learning and

Carbone and Duffy exploring peer-to-peer social learning. Both studies report that subjects have some difficulty with habit-formation specifications as they require that subjects optimally save more early on in their life-cycle (relative to the absence of a habit variable) to adjust for the diminishing effect that habits have on utility over the lifecycle, and consistent with earlier studies (without habit), consumers typically undersave early on in their life-cycle. Brown and others find that information on the life-cycle consumption/savings choices made by prior experienced generations of subjects (intergenerational learning) improves the performance of subsequent generations of subjects (in terms of closeness to the optimal path). However Carbone and Duffy report that social information on the contemporary consumption/savings choices of similarly situated peers (peer-to-peer learning) does not improve performance in the model with (or without) habit in the utility function.

Future experimental research on dynamic, intertemporal consumption/savings plans might explore the impact of other realistic but currently missing features, such as mortality risk, an active borrowing and lending market among agents of different ages, consumption/leisure trade-offs, and the consequences of retirement and social security systems.

2.2 Exponential Discounting and Infinite Horizons

It is common in macroeconomic models to assume infinite horizons, as the representative household is typically viewed as a dynasty, with an operational bequest motive linking one generation with the next. Of course, *infinite* horizons are not operational in the laboratory, but *indefinite* horizons are. As we have seen, in experimental studies, these have often been implemented by having a constant probability δ that a sequence of decision rounds continues with another round.[12] Theoretically this practice should induce both exponential discounting of future payoffs at rate δ per round as well as the stationarity associated with an infinite horizon, in the sense that, for any round reached, the expected number of future rounds to be played is always $\delta + \delta^2 + \delta^3 + \ldots$, or, in the limit $\delta/(1-\delta)$. Empirically, there is laboratory evidence that suggests that probabilistic continuation does affect subjects' perceptions of short-run versus long-run incentives as predicted by theory. For instance, Dal Bó (2005) reports lower cooperation for finite-duration experiments in comparison to indefinite-duration experiments having the same expected length. In particular, Dal Bó reports that aggregate cooperation rates are positively correlated with the continuation probability implemented.

To better induce discounting at rate δ, it seems desirable to have subjects participate in *several* indefinitely repeated sequences of rounds within a given session—as opposed to a single indefinitely repeated sequence—as the former practice provides subjects with the experience that a sequence ends and thus a better sense of the intertemporal rate of discount they should apply to payoffs. A further good practice is to make transparent the randomization device for determining whether an indefinite sequence continues or not, for example, by letting the subjects themselves roll a die at the end of each round using a rolling cup. A difficult issue is the possibility that an indefinite sequence continues beyond the scheduled time of an experimental session. One approach to dealing with this problem is to recruit subjects for a longer period of time than is likely necessary, say, several hours, and inform them that a number of indefinitely repeated sequences of rounds will be played for a set amount of time—say for one hour following the reading of instructions. Subjects would be further instructed at the outset of the session that after that set amount of time had passed, the indefinite sequence of rounds currently in play would be the last indefinite sequence of the experimental session. In the event that this last indefinite sequence continued beyond the long period scheduled for the session, subjects would be instructed that they would have to return at a later date and time that was convenient for everyone to complete that final indefinite sequence.

In practice, as we have seen, some researchers feel more comfortable working with finite-horizon models. However, replacing an infinite horizon with a finite horizon may not be innocuous; such a change may greatly alter predicted behavior relative to the infinite-horizon case. For instance, the finite-horizon life-cycle model of the consumption-savings decision greatly increases the extent of the precautionary savings motive relative to the infinite-horizon case. Other researchers have chosen not to tell subjects when a sequence of decision rounds is to end (e.g., Offerman et al. 2001), or to exclude data from the end rounds (e.g., Ule et al. 2009) as a means of gathering data from an approximately infinite horizon. A difficulty with that practice is that the experimenter loses control of subjects' expectations regarding the likely continuation of a sequence of decisions and appropriate discounting of payoffs. This can be a problem if, for instance, the existence of equilibria depend on the discount factor being sufficiently high. Yet another approach is to exponentially discount the payoffs that subjects receive in each round but at some point in the session switch over to a stochastic termination rule (e.g., Feinberg and Husted 1993). A problem with this approach is that it does not implement the *stationarity* associated with an infinite horizon.

2.3 *Exponential or Hyperbolic Discounting?*

Recently, there has been a revival of interest in time-inconsistent preferences with regard to consumption-savings decisions, where exponential discounting is replaced by a quasi-hyperbolic form so that the representative agent is viewed as maximizing

$$u(c_t) + \beta \sum_{i=1}^{T} \delta^i u(c_{t+i})$$

where $\delta \in (0, 1)$ is a discount factor and the parameter $\beta \leq 1$ characterizes the agent's *bias for the present* (exponential discounting has $\beta = 1$).[13] Agents who discount hyperbolically ($\beta < 1$) rather than exponentially may exhibit time-inconsistent behavior (self-control problems) in that they systematically prefer to reverse earlier decisions, for example, regarding how much they have saved. Thus, a possible explanation for the departures from optimal consumption paths noted before in experimental studies of intertemporal decision making may be that subjects have such present-biased preferences. Indeed, Laibson (1997), O'Donoghue and Rabin (1999), and several others have shown that consumers with such preferences save less than exponential consumers.

Although time-inconsistent preferences have been documented in numerous psychological studies (see, e.g., Frederick, Loewenstein, and O'Donoghue (2002) for a survey) the methodology used has often consisted of showing inconsistencies in hypothetical (i.e., unpaid) money-time choices (e.g., Thaler 1981). For example, subjects are asked whether they would prefer $\$D$ now or $\$D(1+r)$ t periods from now, where variations in both r and t are used to infer individual rates of time preference. Recently, nonhypothetical (i.e., paid) money-time choice experiments have been conducted that more carefully respect the time dimension of the trade-off (e.g., Coller, Harrison, and Rutström (2005); Benhabib, Bisin, and Schotter (2010). These studies cast doubt on the notion that discounting is consistent with either exponential or quasi-hyperbolic models of discounting. For instance, Benhabib and others report that discount rates appear to vary with both the time delay from the present and the amount of future rewards in contrast to exponential discounting. However, Coller and others show that in choices between money rewards to be received only in the future, for example, seven days from now versus thirty days from now, variations in the time delay between such future rewards do not appear to affect discount rates, which is consistent with both exponential and quasi-hyperbolic discounting but inconsistent with continuous hyperbolic discounting. Consistent with quasi-hyperbolic discounting both studies find that a small fixed premium attached to immediate versus delayed rewards, can reconcile much of the variation in discount rates between the present and the future and between different future rewards. However, this small fixed premium does not appear to vary with the amount of future rewards (Benhabib et al.) and may simply reflect transaction/credibility costs associated with receiving delayed rewards (Coller et al.), making it difficult to conclude definitively in favor of the quasi-hyperbolic model.

Anderson et al. (2008) make a strong case that time preferences cannot be elicited apart from risk preferences. Prior studies on time discounting all presume that subjects have risk-neutral preferences. However, if subjects have risk-averse preferences (concave utility functions) as is typically the case, the implied discount rates from the binary time-preference choices will be *lower* than under the presumption of risk neutrality (linear utility functions). Indeed, Anderson et al. (2008) elicit joint time and risk preferences by having each subject complete sequences of binary lottery choices (of the Holt and Laury (2002) variety) that are designed to elicit risk preferences as well as sequences of binary time-preference choices that are designed to elicit their discount rates (similar to those in the Coller et al. study). They find that once the risk aversion of individual subjects is taken into account, the implied discount rates are much lower than under the assumption of risk-neutral preferences. This finding holds regardless of whether discounting is specified to be exponential or quasi-hyperbolic or some mixture.

Of course, one must use caution in extrapolating from experimental findings on intertemporal decision making to the intertemporal choices made by the representative household, firm, government agencies, or institutions in the macroeconomy. Internal, unaccounted-for factors may bias intertemporal decision making in ways that experimental evidence cannot easily address; for example, election cycles or other seasonal factors may influence decision making in ways that would be difficult to capture in a laboratory setting.

2.4 Expectation Formation

In modern, self-referential macroeconomic models, expectations of future endogenous variables play a critical role in the determination of the current values of those endogenous variables; that is, beliefs affect outcomes, which in turn affect beliefs, which affect outcomes, and so on. Since Lucas (1972) it has become standard practice to assume that agents' expectations are *rational* in the sense of Muth (1961), and indeed most models are "closed" under the rational expectations assumption. The use of rational expectations to close self-referential models means that econometric tests of these models using field data are joint tests of the model and the rational expectations assumption, confounding the issue of whether the expectational assumption or other aspects of the model are at fault if the econometric evidence is at odds with theoretical predictions. While many tests of rational expectations have been conducted using survey data, (e.g. Frankel and Froot 1987), these tests are beset by problems of interpretation, for example, due to uncontrolled variations in underlying fundamental factors, or to the limited incentives of forecasters to provide accurate forecasts, or to disagreement about the true underlying model or data-generating process. By contrast, in the lab it is possible to exert more control over such confounding factors, to know for certain the true data-generating process, and to implement the self-referential aspect of macroeconomic models.

Early experimental tests of rational expectations involved analyses of subjects' forecasts of exogenous, stochastic processes for prices, severing the critical self-referential aspect of macroeconomic models but controlling for the potentially confounding effects of changes in fundamental factors (e.g., Schmalensee 1976; Dwyer et al. (1993). Later experimental tests involved elicitation of price forecasts from subjects who were simultaneously participants in experimental asset markets that were determining the prices being forecast (Williams 1987; Smith, Suchanek, and Williams (1988). As discussed in the prior handbook surveys by Camerer (1995) and Ochs (1995), many (though not all) of these papers found little support for rational expectations in that forecast errors tended to have nonzero means and were autocorrelated or were correlated with other observables. Further, the path of prices sometimes departed significantly from rational expectations equilibrium. However, most of these experimental studies involve analyses of price forecasts in environments where there is no explicit mechanism by which forecasts determine subsequent outcomes, as is assumed in forward–looking macroeconomic models. Further, some of these experimental tests (e.g., Smith et al.) involved analyses of price forecasts for relatively short periods of time or in empirically nonstationary environments where trading behavior resulted in price bubbles and crashes, providing a particularly challenging test for rational expectations hypothesis.

Marimon and Sunder (1993, 1994) recognized the challenge to subjects of both forecasting prices and then using those forecasts to solve complicated dynamic optimization problems. They pioneered an approach that has come to be known as a *learning-to-forecast* experimental design, another methodology that might be considered "native" to experimental macroeconomics. In their implementation, subjects were asked each period to form inflationary expectations in a stationary overlapping-generations economy. These forecasts were then used as input into a computer program that solved for each individual's optimal, intertemporal consumption/savings decision given that individual's forecast. Finally, via market clearing, the actual price level was determined and therefore the inflation rate. Subjects were rewarded only for the accuracy of their inflation forecasts and not on the basis of their consumption/savings decision, which was, after all, chosen for them by the computer program. Indeed, subjects were not even aware of the underlying overlapping-generations model in which they were operating—instead they were engaged in a simple forecasting game. This learning-to-forecast approach may be contrasted with a "learning-to-optimize" experimental design, wherein subjects are simply called upon to make choice decisions (e.g. consumption/savings) having intertemporal consequences but without elicitation of their forecasts (which are implicit). This is an interesting way of decomposing the problem faced by agents in complex macroeconomic settings so that it does not involve a joint test of rationality in both optimization and expectation formation; indeed, the learning to forecast experimental design has become a workhorse approach in experimental macroeconomics—see Hommes (2011) for a comprehensive survey.

More recently some macroeconomists have come to believe that rational expectations presumes too much knowledge on the part of the agents who reside within these models. For instance, rational expectations presumes common knowledge of rationality. Further, rational expectations agents know with certainty the underlying model, whereas econometricians are often uncertain of data-generating processes and resort to specification tests. Given these strong assumptions, some researchers have chosen to replace rational expectations with some notion of bounded rationality and ask whether boundedly rational agents operating for some length of time in a known, stationary environment might eventually *learn* to possess rational expectations from observation of the relevant time-series data (see, e.g., Sargent (1993, 1999) and Evans and Honkapohja (2001) for surveys of the theoretical literature).

Learning to forecast experiments have played a complementary role to the literature on learning in macroeconomic systems. This literature imagines that agents are boundedly rational in the sense that they do not initially know the model (data-generating process) and behave more as econometricians, using possibly misspecified model specifications for their forecasting rules, which they update in real time as new data become available. In addition to the work of Marimon and Sunder (1993, 1994), this real-time, adaptive expectations approach has been explored experimentally using the learning to forecast design by Bernasconi, Kirchkamp, and Paruolo (2006), Hey (1994), Van Huyck Cook, and Battalio (1994), Kelley and Friedman (2002), Hommes and others (2005, 2007), Heemeijer and others (2009), and Bao and others (2012), Bao, Duffy, and Hommes (2013). The use of the learning to forecast methodology has become particularly important in assessing policy predictions using the expectations-based New Keynesian model of the monetary-transmission mechanism in experimental studies by Adam (2007), Pfajfar and Zakelj (2015), Assenza and others (2013), and Petersen (2015), as will be discussed in Section 5.3.

Hommes and others (2007) provide a good representative example of this literature. They consider expectation formation by groups of six subjects operating for a long time (in the laboratory sense)—fifty periods—in the simplest dynamic and self-referential model, the cobweb model.[14] In each of the fifty periods, all six subjects are asked to supply a one-step-ahead forecast of the price that will prevail at time t, $p^e_{i,t}$, using all available past price data through time $t-1$; the forecast is restricted to lie in the interval (0, 10). These price forecasts are automatically converted into supply of the single good via a supply function $s(p^e_{i,t};\ \lambda)$, which is increasing in $p^e_{i,t}$ and has common parameter λ governing the nonlinearity of the supply function. Demand is exogenous and given by a linear function $D(p_t)$. The unique equilibrium price p^* is thus given by

$$p^*_t = D^{-1}\left(\sum_{i=1}^{6} s(p^e_{i,t})\right)$$

that is, it is completely determined by subjects' price forecasts. However, Hommes and others add a small shock to exogenous demand, which implies that prices should evolve according to $p_t = p^*_t + \epsilon$, where $\epsilon \sim N(0, \sigma^2_\epsilon)$. Thus, under rational expectations, all forecasters should forecast the same price, p^*. In the new learning view of rational expectations, it is sufficient that agents have access to the entire past history of prices for learning of the rational-expectations solution to take place. Consistent with this view, Hommes and others do not inform subjects of the market clearing process by which prices are determined. Instead, subjects are simply engaged in forming accurate price forecasts and individual payoffs are a linearly decreasing function of the quadratic loss $(p_t - p^e_{i,t})^2$. The main treatment variable consists of variation in the supply function parameter λ, which affects the stability of the cobweb model under the assumption

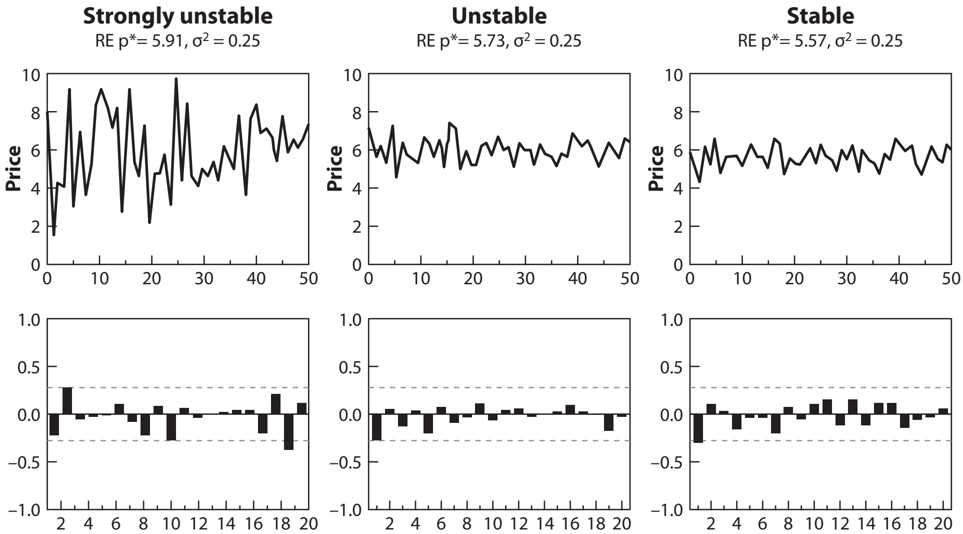

Figure 1.2: Actual prices (top) and autocorrelations (bottom) from three representative sessions of the three treatments of Hommes et al. (2007): strongly unstable, unstable, and stable equilibrium under naïve expectations.

of *naive* expectations (following the classic analysis of Ezekiel (1938)). The authors consider three values for λ, for which the equilibrium is stable, unstable, or strongly unstable under naive expectations.[15] Their assessment of the validity of the rational-expectations assumption is based on whether market prices are biased (looking at the mean), whether price fluctuations exhibit excess volatility (looking at the variance), and whether realized prices are predictable (looking at the autocorrelations).

Figure 1.2 shows a representative sample of prices and the autocorrelation of these prices from the three representative groups operating in the three different treatment conditions. This figure reveals the main finding of the study, which is that in all three treatments, the mean price forecast is not significantly different from the rational expectations value, though the variance is significantly greater (there is excess volatility) from the rational expectations value, $\sigma_\epsilon^2 = 0.25$, in the unstable and strongly unstable cases. Even more interesting is the finding that the autocorrelations are not significantly different from zero (5% bounds are shown in the figures) and there is no predictable structure to these autocorrelations. The latter finding suggests that subjects are not behaving in an irrational manner in the sense that there is no unexploited opportunities for improving price predictions. This finding is somewhat remarkable given the limited information subjects had regarding the model generating the data, though coordination on the rational expectations equilibrium was likely helped by having a unique equilibrium and a limited price range (0, 10).

Adam (2007) uses the learning to forecast methodology in the context of the two-equation, multivariate New Keynesian "sticky price" model that is a current workhorse of monetary policy analysis (e.g., Woodford 2003).[16] In a linearized version of that model, inflation, π_t, and output, y_t, are determined by the system of expectational difference equations,

$$\begin{pmatrix} \pi_t \\ y_t \end{pmatrix} = a_0 + a_1 y_{t-1} + B \begin{pmatrix} \pi_t^e \\ \pi_{t+1}^e \end{pmatrix} + c v_t$$

where a_0, a_1, B, and c are conformable vectors and matrices, π_t^e and π_{t+1}^e are the one- and two-step-ahead forecasts of future inflation using information available through time $t-1$, and v_t is a mean zero real monetary shock. Like Hommes and others, Adam provides information on all past realizations of π and y through period $t-1$ and asks a group of five subjects to provide one- and two-step-ahead forecasts of inflation, π_t^e and π_{t+1}^e, repeatedly for forty-five to fifty-five periods. The average forecasts each period are used in the model above to determine π_t and y_t. Subjects earn payoffs based on forecast accuracy alone and are uninformed regarding the underlying process generating data on π_t and y_t.

The rational expectation solution is of the form

$$y_t = y + v_t$$

$$\pi_t = \left(\frac{\pi}{y}\right) y_{t-1}$$

where y and π represent steady-state values. Inflation lags output by one period due to predetermined (sticky) prices, and output deviates from its steady state only due to real monetary shocks. Thus a rational forecast model for π_t should condition on y_{t-1}, that is, $\pi_t = \alpha_y + \beta_y y_{t-1}$. Of course, since subjects are given time-series data on both y and π, Adam imagines that subjects might alternatively use a simple (but misspecified) autoregressive forecast model of the form $\pi_t = \alpha_\pi + \beta_\pi \pi_{t-1}$. Thus, the issue being tested here is not simply one of whether agents can learn to form rational expectations of future inflation but more importantly whether subjects, like econometricians, can find the correct specification of the reduced-form model they should use to form those rational expectations. Perhaps not surprisingly, the evidence on the latter question is somewhat mixed. Adam finds that in most of the experimental sessions, subjects forecast using the autoregressive inflation model and do not condition their forecasts on lagged output. However, he also shows that such behavior can result in a stationary, "restricted-perceptions" equilibrium that is optimal in the sense that autoregressive inflation forecasts outperform those that condition on lagged output. Adams further notes that this miss-specification in agents' forecasts provides a further source of inflation and output persistence in addition to that implied by the model's assumption of sticky price adjustment, a finding that has been elaborated upon by Davis and Korenock (2011).

Bao and others (2013) study learning behavior in a cobweb model with a setup similar to that of Hommes and others (2007). However, they compare the performance of the learning-to-forecast experimental design with the alternative "learning-to-optimize" design, where subjects in the role of suppliers must directly choose the *quantity*, q_t^i, of the good they wish to bring to the market in period t. In the latter case, the quantity of the six agents is simply summed up to give aggregate supply. Market clearing using the exogenous market demand yields the market price, p_t. Subjects in this learning–to–optimize design are paid on the basis of their profit, $p_t q_t^i - c(q_t^i)$, where $c(\cdot)$ is a known convex cost function. Bao and others have two further treatments: one in which subjects are asked to both form price forecasts and choose supply decisions and a second in which two subject teams are formed, with one team member performing the forecasting task that the other team member could use to determine the quantity task. In the latter two treatments, subjects are paid an equal weighted average of the payoffs from the forecasting and profit-maximizing tasks. Bao and others report that convergence to the rational-expectations equilibrium (REE) is fastest in the learning-to-forecast design and slowest and highly variable in the treatment where individual subjects must both forecast and choose quantity decisions. Dividing up the two tasks among team members greatly improves performance. These findings indicate that learning-to-forecast designs should be regarded as an upper bound on the speed and efficiency with which agents may learn a REE and that it may be more useful to think of the representative household or firm as a team of specialized actors.

A second approach to boundedly rational expectation formation in macroeconomics takes into account the strategic uncertainties that can arise from interactions among heterogeneous agents. This approach is sometimes referred to as *step-level* reasoning and was motivated by Keynes's (1936) famous comparison of financial market investors' expectations to newspaper beauty contests of that era, in which participants had to select the 6 prettiest faces from 100 photographs. The winner of the contest was the person whose choices were closest to the average choices of all competitors. Keynes (1936, 156) noted that "each competitor has to pick, not those faces which he himself finds prettiest but those he thinks likeliest to catch the fancy of other competitors, all of whom are looking at the problem from the same point of view." Keynes went on to observe that individuals might form expectations not just of average opinion, but might also consider what average opinion expects average opinion will be, and he further speculated that there might be some who practiced still "higher degrees" of reasoning. These observations concerning expectation formation were tested experimentally by Nagel (1995) in a game developed by Moulin (1986) that has since come to be termed the "beauty contest" game in honor of Keynes's analogy.

In Nagel's design, a group of $N = 15 - 18$ subjects are each asked to "guess"—simultaneously and without communication—a real number in the closed interval [0, 100]. They are further instructed that the person(s) whose guess is closest in absolute value to a known parameter p times the mean of all submitted numbers is the winner of a large cash prize, while all other participants receive nothing. Nagel's baseline experiment involves setting $p < 1$, for example, $p = \frac{2}{3}$. That game is straightforward to analyze: each player i wants to guess a number $x_i = p\bar{x}$, where $\bar{x}$ is the mean of all submitted numbers. Given this objective, in any rational expectations equilibrium we must have that $x_i = \bar{x}$ for all i. If $p < 1$, the only rational expectations solution is $x_i = \bar{x} = 0$, that is, all N players guess *zero*.[17] To map this game into Keynes's (1936) example requires setting $p = 1$, in which case any number in [0, 100] is a rational expectations equilibrium; the choice of $p < 1$ yields not only a unique equilibrium prediction but interesting insights regarding the extent of individual's higher degrees of reasoning.[18]

Nagel's experimental findings from three sessions of the $p = \frac{1}{2}$-mean game are shown in Figure 1.3, which reports the relative frequencies of number choices in the interval [0, 100].[19] Notice first that the equilibrium prediction of 0 is never chosen. Second, there are large spikes in neighborhoods of the numbers 50, 25, and 12.5. A choice of 50 implies an expected mean of 100 in the $p = \frac{1}{2}$ game and is thus *barely* rational—these players exhibit the lowest level of reasoning, which is often termed *step*, or *level*, 0. The somewhat more sophisticated level 1 types expect a mean of 50 and guess numbers that are $\frac{1}{2}$ of their expectation around 25, while level 2 types are a step further ahead, anticipating a mean of 25 and thus guessing numbers around 12 or 13. A robust finding is that depths of reasoning in excess of level 2 are rarely observed; the winner of the beauty contest is typically a level-2 type. With repetition, subjects in these beauty contest games do eventually converge upon the unique rational expectations equilibrium prediction (0 in this case), but each

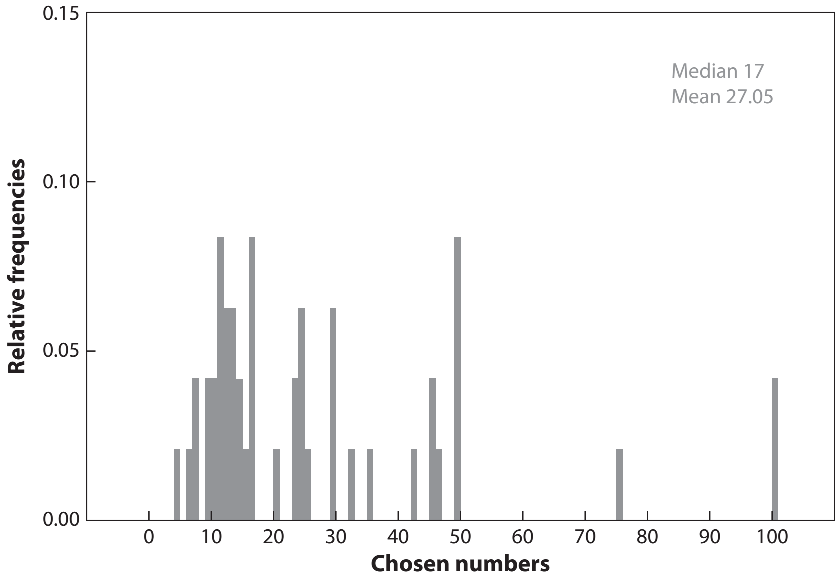

Figure 1.3: Relative frequencies of numbers in the interval [0, 100] chosen in Nagel's $\frac{1}{2}$-mean game (beauty contest) Source: Nagel (1995).

individual's process of expectation revision over time typically follows the same level of reasoning they exhibited in the first round played, for example, level $k = 1$ or 2 adjustment in each repetition. This experiment, which has now been replicated many times (see, e.g., Duffy and Nagel 1997; Ho, Camerer, and Weigelt 1998), reveals that in multiagent economies where all agents know the model, the common-knowledge-of-rationality assumption implicit in the rational expectations hypothesis may not hold. It further suggests that decision costs or cognitive constraints may lead individuals to adopt heuristic rules of thumb that result in predictable step-levels of belief revision, that is, *systematic* forecast errors. That convergence to equilibrium does obtain in the limit is reassuring but suggests that rational expectations might be best viewed as a long-run phenomenon.

Summing up, we have seen some ways in which three microlevel assumptions that are mainstays of macroeconomic modeling—intertemporal optimization, time-consistent preferences/exponential discounting, and the rationality of expectations—have been tested in the laboratory, primarily in individual decision-making experiments. The evidence to date suggests that human subject behavior is often at odds with the standard micro-assumptions of macroeconomic models. The behavior of subjects appears to be closest to microassumptions, for example, intertemporal optimization, when subjects learn from one another or gather information on prices through participation in markets. The rational expectations model appears to be most reasonable in simple, univariate models (e.g., the cobweb model) as opposed to the more commonly used multivariate models. Hopefully, these and other experimental findings will lead to a reconsideration of the manner in which macroeconomic modelers characterize the behavior of their "representative" agents, though so far, there is not much evidence that such a change is imminent.

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 $\overline{k}_h$ representing the Pareto efficient equilibrium. The dynamics of the system (under perfect foresight) are such that for $k \in (0, k^*)$, $\overline{k}_l$ is an attractor, whereas for $k \geq k^*$, $\overline{k}_h$ 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 $\overline{k}_h$ 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

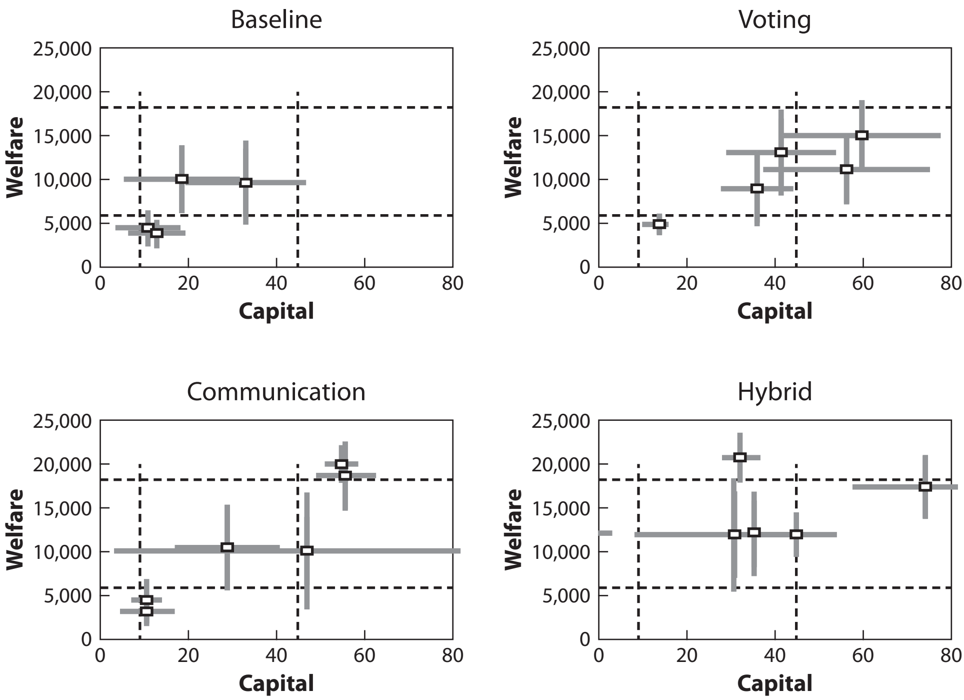

Figure 1.4: Asymptotic estimates of aggregate welfare (vertical axis) and capital (horizontal axis) for each session (square) of the four treatments of Capra et al. (2009).

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.

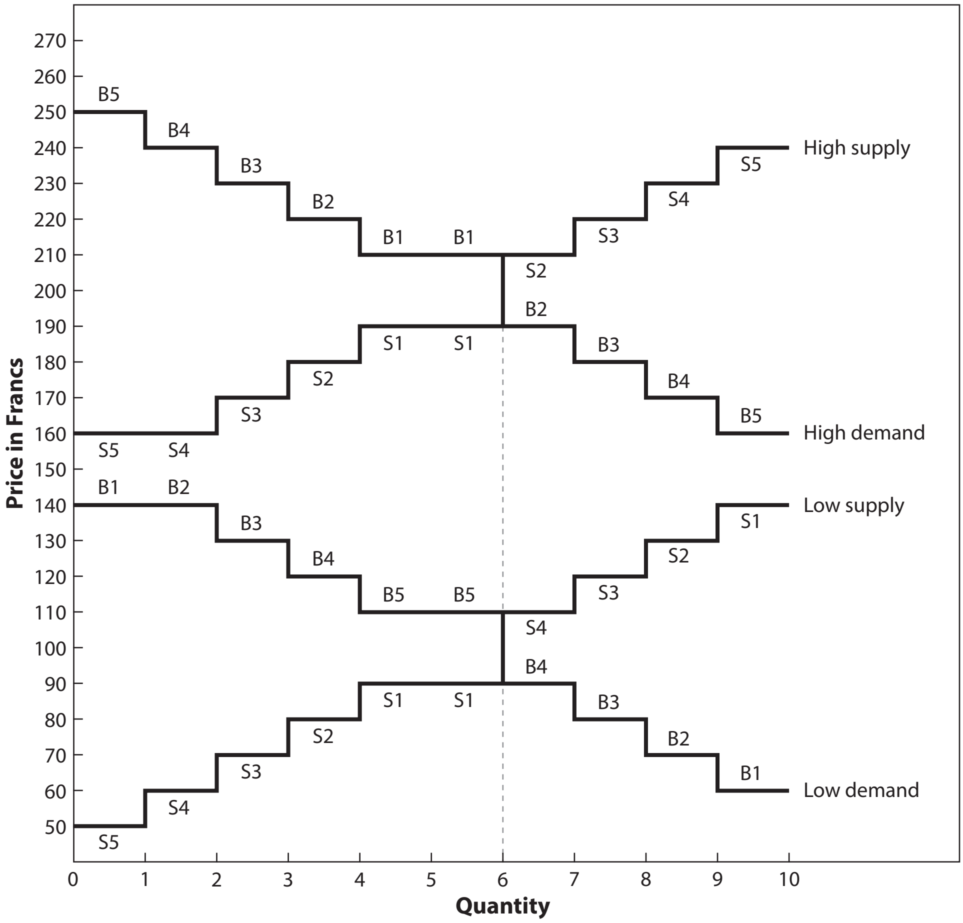

Figure 1.5: Induced high and low demand and supply in Duffy and Fisher (2005).

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