

“Lucas” In The Laboratory*

ELENA ASPAROUHOVA, PETER BOSSAERTS,

NILANJAN ROY, and WILLIAM ZAME

ABSTRACT

We study the Lucas asset pricing model in a controlled setting. Participants trade two long-lived securities in a continuous open-book system. The experimental design emulates the stationary, infinite-horizon setting of the model and incentivizes participants to smooth consumption across periods. Consistent with the model, prices align with consumption betas and co-move with aggregate dividends, particularly so when risk premia are higher. Trading significantly increases consumption smoothing compared to autarky. Nevertheless, as in field markets, prices are excessively volatile. The noise corrupts traditional GMM tests. Choices display substantial heterogeneity, with no subject representative for pricing.

For over 30 years, the Lucas intertemporal asset pricing model (Lucas (1978)) and its extensions and variations have served as the basic platform for research on dynamic asset pricing and business cycles. At the cross-sectional level, the Lucas model predicts that only aggregate consumption risk is priced.¹ At the time-series level, the Lucas model predicts that the level and volatility of asset prices are correlated with the level and volatility of aggregate consumption.

In this paper we study multiperiod, multisecurity asset pricing in a controlled setting and evaluate the outcomes against the predictions of the Lucas model. We find that consumption beta ranks security prices in the cross section, as predicted by the model, but not security returns over time. We also find, as the model predicts, that fundamentals – aggregate consumption – drive changes in prices over time. Because of our parametrization, this implies significant predictability, in violation of simple accounts of the efficient markets hypothesis (EMH). Nevertheless, as in the field, the theory is rejected because prices are excessively volatile: fundamentals explain at most 18% of price changes, while the theory predicts that fundamentals should explain 100% of price changes.

Since incentives are controlled and all choices are observed in the lab, we are able to study welfare. Across participant types (distinguished by initial holdings and period-income fluctuations), we discover substantial Pareto improvements from autarky, to the extent that consumption shares become statistically insignificantly different from constant, both over time and across states.² Thus, excess volatility does not appear to prevent improvements in welfare. However, we document considerable individual heterogeneity, so that no individual

choice can be considered representative for pricing. Closer inspection of choices suggests that excess volatility is consistent with decision making based on subjects' (only mildly incorrect) prediction that prices do not change with fundamentals. The simulation of an economy populated with agents endowed with such beliefs leads to market outcomes that closely resemble the outcomes of our experiment.

Finally, we study instrumentation in generalized method of moments (GMM) tests of the asset return–consumption restrictions of the Lucas model and find that correct inference requires observation of the true underlying state of the economy. Unless the true state is used as an instrument (one that is hard to envision as available to the econometrician using field data), the GMM tests with traditional instruments (lagged returns, consumption) generate incorrect parameter estimates and suggest misleading conclusions about the validity of the model.

From a methodological point of view, our novel experimental design addresses several difficulties that the researcher faces in conducting laboratory tests of the Lucas model, such as the need to generate a stationary, infinite-horizon setting and to make period-consumption perishable. Prices in our experiment do not “bubble,” a finding that is in contrast with the vast literature on laboratory bubbles and that we attribute to our design featuring incessant incentives to trade. This echoes the opinion of Crockett and Duffy (2013).

We proceed as follows. Because experimental tests of asset pricing theory are still rare, in Section I we provide extensive motivation and in Section II we discuss the challenges one

encounters when attempting to test the Lucas model in a controlled setting. In Section III we present the Lucas model within a stylized version of the framework that we created in the laboratory. Section IV provides details of the experimental design. Results are discussed in Section V. Section VI uses the data from the laboratory to investigate the statistical properties of the GMM tests with which historical data from the field have traditionally been analyzed. Section VII discusses potential causes of the excessive volatility of asset prices observed in our laboratory markets. Section VIII concludes. Appendix A discusses the theory in greater detail. Appendix B presents an illustrative numerical example. Finally, Appendix C reproduces the instructions and timeline provided to the laboratory participants.

I. Why Experimental Tests of The Lucas Model?

Controlled experimentation with markets is not standard methodology in empirical finance, and thus we first must address why we think our exercise has value. We believe that experimentation itself should not be an issue. Controlled experiments are the foundation of science. That experimentation is still rare in finance is likely due to the difficulty researchers face in designing experiments that are informative. We argue that experiments can be “informative” even if the experimental setting does not match exactly the “real world.” The goal of experiments is, in the first place, to test the veracity of theory in a setting where confounding factors are eliminated as much as possible.

The Lucas model has been tested exclusively³ on historical data from the field using statistical analysis⁴ or calibration⁵ of its core equilibrium restrictions, the stochastic Euler equations.⁶ It is fair to say that the general conclusion from these tests is that the model fails. One may wonder why we would bother testing in the laboratory something that is “obviously wrong” in the field. We argue that it is precisely *because* the model fails on historical data that one should test whether it is true using controlled experiments. The fact that the stochastic Euler equations do not fit a given set of historical data from the field does not make the model invalid. Observation in the field is inevitably incomplete. For instance, one may be using the wrong measurements (e.g., consumption of durable goods is left out or modeled incorrectly⁷), there may be forces at play that the theory abstracts from and as a result are ignored in the empirical analysis (e.g., transaction costs), or the theory may not apply because an important assumption is violated (e.g., allocations are not Pareto efficient, and hence a representative agent does not exist).

Our experimental design represents a realistic setting, yet with minimal complexity. Real people trade for real money in real markets, and their task captures the two key goals of trading in the Lucas model, namely, diversification (across risky securities) and smoothing (of consumption over time). Our design accommodates agent heterogeneity (human subjects exhibit differences in, for example, attitudes towards risk), and imposes endowments that are nonstationary in order to induce trade beyond the first period.

At the same time, our experiment does not provide a faithful replication of the field. One deviation in particular is that the experiment uses a design with dividends that are

stationary in levels and not in growth. In contrast, empirical tests of the Lucas model on historical data from the field build on an extension of the Lucas model that is stationary in consumption growth rather than in consumption levels (Mehra and Prescott (1985)). For obvious practical reasons, we stay with the original version. Indeed, incentives would be distorted if payment levels were to depend on how long a session lasts (with stationary consumption growth, dividends increase over time, which means that subjects are paid more in sessions with more periods). Most importantly, however, because of our assumption of stationarity in levels, the main conclusions we draw do not depend on particular functional forms (of preferences) even if one allows, as we do, for heterogeneity across agents.

There is one important difference between the stationary-in-levels and stationary-in-growth models. Stationary-in-levels models readily generate equilibrium phenomena such as countercyclical equity premia without having to appeal to countercyclical risk aversion (Routledge and Zin (2011)). Assuming standard preferences such as power utility, when the economy is stationary in levels, wealth is higher in a good state than in a bad state (states are defined as levels of wealth). As a result, the curvature of the utility function is lower in a high state than in a low state, so risk premia are lower in a high state than in a low state. One thereby obtains countercyclical equity premia. When the economy is stationary in growth, however, the state can be good (consumption growth is high leading into the period) or bad (consumption growth is low) with probabilities that are independent of wealth levels. (Notice that here states are defined in terms of growth, not wealth levels.) Since utility curvature depends only on wealth levels, and the state of the economy is unrelated to

wealth levels, one cannot have countercyclical equity premia. To nevertheless obtain countercyclical equity premia, one could introduce a reference point that is low when the state is good and high when the state is bad (Campbell and Cochrane (1999)). This way, the curvature of the utility function depends not only on wealth, but also explicitly on the state. Countercyclical equity premia emerge naturally.

While the goal of experiments is not to replicate the field, experimental results become potentially more informative if mismatches between theory and (experimental) data are not unlike the anomalies one observes in the field. The excess volatility that we observe in our experiments happens to be one of the dominant puzzles about historical financial markets (Shiller (1981)). Interestingly, excessive volatility emerges in our experiments despite the fact that the arguments used to explain the phenomenon in the field do not apply to our setting. These include: durability of consumption goods (Dunn and Singleton (1986)), preference for early resolution of uncertainty (Epstein and Zin (1991)), collateral use of certain consumption goods (Lustig and Nieuwerburgh (2005)), and long-run dependencies in the statistical properties of consumption flows (Bansal and Yaron (2004)).

In the laboratory, inspection of an anomaly may reveal the blind spots of the theory, and therefore help identify why the theory fails to explain the field. Prices in the Lucas model are determined solely by (“are measurable in”) the fundamental risk in the economy, and price risk is equivalent to fundamental risk. In our experiments, there is only one source of fundamental risk, namely, aggregate consumption.⁸ Yet a large fraction of the price changes we observe is independent of changes in fundamentals. The existence of this apparent residual

risk, which is distinct from fundamental risk, is inconsistent with the Lucas model. Within the model any such residual risk is eliminated by agents' perfect knowledge of how prices change with fundamentals, that is, by their perfect foresight. We not only observe this residual risk, which we refer to as *residual price forecasting risk*,⁹ but we document that it is large in magnitude: fundamentals explain at most 18% of the variability of securities prices, so residual price forecasting risk accounts for the remaining 82%.¹⁰

In retrospect, it perhaps should have been obvious that imperfect foresight is the Achilles heel of the Lucas model. But in fact it did not seem so obvious. Starting with Plott and Sunder (1988), evidence from financial market experiments with heterogeneous information provides overwhelming support for the view that markets can produce perfect foresight equilibria, to the extent that markets are promoted as forecasting devices (Arrow et al. (2008)). We are beginning to understand the aspects of human nature that facilitate this remarkable ability to acquire perfect foresight (Bruguier, Quartz, and Bossaerts (2010)). There, however, the setting is one in which differentially informed participants have to extract insider information from order and trade flows. In the context we consider, equally informed participants have to build, from the same flows, correct expectations of how future clearing prices relate to fundamentals.¹¹

As we shall see, the design we use to test the Lucas model requires that participants follow quite sophisticated hedging strategies. One may conjecture that our subjects (students at Caltech, UCLA, and the University of Utah) lack sufficient financial sophistication to conceive of, let alone execute, these strategies. However, plenty of examples in the experimental

economics, psychology, and computational neuroscience literatures show that humans are capable of quite complicated behavior towards uncertainty. We already alluded to subjects' remarkable ability to trade on other participants' information (Plott and Sunder (1988)). In a two-period asset market setting, Bossaerts, Meloso and Zame (2013) show that subjects can resort to complex hedging strategies to improve performance.

Controlled experiments can also shed light on the plausibility of important features of a theory in the presence of a more compelling alternative. Such is the case with the Lucas model, where prices are tightly linked to fundamentals – to the extent that fundamentals are predictable, prices must be predictable as well. In contrast, original accounts of the EMH state that prices must not be predictable, but rather should be a martingale, except for drift as compensation for risk (Samuelson (1973), Malkiel (1999)). The idea behind the EMH is that investors trade to exploit the predictability and in the process eliminate it. In historical data from the field, many exceptions have been discovered. Securities prices can be predicted with, for example, the dividend yield (although not everyone agrees; see Goyal and Welch (2003) for a critique), securities prices exhibit cyclicalities (Lo and MacKinlay (1988)), and price evolution can be predicted for several months after a specific event such as earnings surprise (Bernard and Thomas (1989)). Interestingly, these violations of EMH tend to be explained, not as confirmation of the Lucas model, but in terms of behavioral finance, which is to say, in terms of cognitive biases in investor decision making (DeBondt and Thaler (1985), Daniel, Hirshleifer, and Subrahmanyam (1998), De Long et al. (1990)). Lucas' proposition is that securities prices *may be predictable* even in properly functioning

markets that do not display any behavioral anomalies. Our experiments demonstrate that this is not only a theoretical possibility, but also eminently relevant to real financial markets. In the experiments, prices move with the aggregate dividend (although in an extremely noisy way). Because aggregate dividends are predictable in the sense that the aggregate dividend is expected to fall in periods with a high aggregate dividend, and vice versa, prices are predictable too. Significantly, the aggregate dividend is the *only* variable that predicts price changes in our experiments.

II. Challenges in Designing an Experimental Test of The Lucas Model

The Lucas model makes assumptions – most importantly, stationarity – that are difficult to satisfy, even in a controlled laboratory setting. Moreover, the model lacks many institutional details, in particular, it simply assumes that markets generate Pareto optimal allocations, so that a representative agent exists, and studies the prices that would support these allocations without considering how such allocations might come about. As a result, designing a laboratory instantiation of the Lucas model is a challenging exercise.

Let us start with the last challenge, the existence of a representative agent. Unless agents are identical, which seems hardly more likely in the laboratory than in the field, the representative agent is only an equilibrium construct, and not a testable assumption/prediction.

Fortunately for us, the heterogeneous agent version of the Lucas model yields predictions that are qualitatively no different than the predictions of the representative agent model (although they arise in a different way) and are testable in the laboratory environment. Pareto optimality plays a central role here. In the representative agent model, Pareto optimality is tautological – there is, after all, only one agent. In the heterogeneous agent model, a representative agent can be constructed, but only if it is assumed that the result of trade is a Pareto optimal allocation – which is not guaranteed – and the particular representative agent that is constructed depends on the particular Pareto optimal allocation that obtains through trade.

For the market outcome to be Pareto optimal would seem to require that the market reach a Walrasian equilibrium. In turn, this would seem to require a complete set of markets, which is an impossibility in an infinite-horizon economy with uncertainty. However, it is enough that markets be *dynamically complete*, which can be the case even with a few assets provided that these assets are long-lived and can be traded frequently (Duffie and Huang (1985)), participants are able to *perfectly foresee* future prices (Radner (1972)), and agents can employ investment strategies that exhibit the *hedging* features that are at the core of the modern theory of both derivatives pricing (Black and Scholes (1973), Merton (1973b)) and dynamic asset pricing (Merton (1973a)).

The second challenge is that agents must learn a great deal. To minimize the difficulty of learning, and consistent with the Lucas model, agents in our experimental economy are told the *exogenous* uncertainty – the dividend process. They still must learn/forecast the

endogenous uncertainty – the price process. As we shall see, this presents agents with a very difficult problem indeed.

In addition to these two challenges, three other aspects of the Lucas model need to be addressed before one can test it in the laboratory. The model assumes that (i) the time horizon is infinite and that agents discount the future, (ii) agents prefer to smooth consumption over time, and (iii) the economy is stationary. Meeting these challenges requires a novel experimental design. We deal with the infinite horizon as in Camerer and Weigelt (1996) by introducing a random ending time determined by a constant termination probability.¹² More precisely, the experimental market unfolds period by period, and each such period can be terminal with the said probability. Termination uncertainty resolves at the conclusion of each period, after subjects have established their securities and cash holdings for that period. Upon termination, all securities expire worthless.

We provide an incentive for participants to smooth consumption by emulating perishability of consumption in each period: at the end of every nonterminal period, holdings of cash (the consumption good) disappear; only cash held at the end of the randomly determined terminal period is credited to participants' final payout accounts (and hence "consumed").

Stationarity of the laboratory economy might seem evident, given the stationarity of the dividend process. However, stationarity of the termination probability presents a severe difficulty. If an experimental session lasts for, say, two hours and each period within that session lasts for four minutes, it would be quite easy for participants to believe that the

termination probability is the same at the end of the first period, when four minutes have elapsed, as it is at the end of the second period, when eight minutes have elapsed – but it may be quite hard for participants to believe that the termination probability remains the same at the end of the 29th period, when 116 minutes have elapsed and only four minutes remain. In that circumstance, participants would surely believe that the termination probability must be higher. However, if as a result subjects believe that the termination probability is not constant, a random ending time would induce a nonconstant discount factor – and very likely induce different discount factors across subjects, because they may each imagine a different way of ensuring that the last period occurs before the session ends. To address this challenge, we introduce a novel termination rule.¹³ We elaborate on this rule in Section IV. Here it suffices to note that it relies on the assumption of stationarity of dividends in levels (and not in growth) and on the small (only two) number of possible dividend realizations.

In concurrent work, Crockett and Duffy (2013) also study an infinite-horizon asset market in the laboratory, but their experimental approach and purpose are different from ours. First, there is no risk in their setting. Our experiment includes risk, and thus allows us to study the interplay between the two core drivers of the Lucas model, namely, risk avoidance (through diversification) and intertemporal consumption smoothing. Second, Crockett and Duffy (2013) induce a preference for consumption smoothing by imposing a schedule of final payments to participants that is nonlinear in period earnings. We take a different approach, and induce preference for consumption smoothing by paying only for one carefully chosen period (the last one), forfeiting payments in all periods that end not being terminal. Third,

the focus in Crockett and Duffy (2013) is on the contrast between pricing in a treatment where there are incentives to trade because of demand for consumption smoothing and pricing in a treatment where there are no such incentives because there is no demand for consumption smoothing. Crockett and Duffy (2013) thus shed light on a long line of experimental work on asset price bubbles, starting with Smith, Suchanek, and Williams (1988), where there are incentives to trade only in initial stages. In contrast, here we are interested in aspects of the Lucas model that have generated controversy in studies of historical asset prices. Of course, for good experimental control, and in line with Crockett and Duffy (2013), our setting is one in which incentives to trade remain present throughout. A final difference concerns the stationarity assumption of the Lucas model. Like us, Crockett and Duffy (2013) uses random termination to induce discounting. But we also need to ensure stationarity, which we obtain through our novel termination protocol.

III. The Lucas Asset Pricing Model in a Setting Amenable To Experimental Tests

We formulate an instantiation of the Lucas asset pricing model that is simple enough to be implemented in the laboratory and yet complex enough to generate a rich set of predictions about prices and allocations, even under very weak assumptions. In particular, we allow for agents with different preferences and endowments (of assets and time-varying income

streams), and we make no assumptions about functional forms, yet we still obtain strong and testable implications for individual consumption choices, trading patterns, and prices.

The design is somewhat complex, but at the outset we should reassure the reader that the task that is required from the subjects is less complex, and is actually rather intuitive. In order for a subject to perform successfully, she does not need to know all the details of the design. It suffices that she appropriately reacts to prices in the marketplace, and that she anticipates future prices reasonably well. At the end of this section, we explain the experiment from the point of view of a subject. (Another way to get a better understanding of the perspective of the subject is to study the Appendix, where we reproduce the instructions.)

A full theoretical analysis of the economy we emulate in the laboratory is delegated to Appendix A. Here, we focus on predictions in the theory that have empirical relevance, that is, that could potentially be observed in our experiments. These concern prices (and returns) in our laboratory markets, as well as subject choices (consumption, trading strategies).

To create an environment suitable for the laboratory setting, we use a formulation that necessarily generates a great deal of trade; in our formulation, Pareto optimality (hence, equilibrium) requires that trading takes place *every period*. This is important in the laboratory setting because subjects do not know the “correct” equilibrium prices (nor do we) and can only learn them through trade, which would seem problematic (to say the least) if trade were to take place infrequently or not at all. We therefore follow Bossaerts and Zame

(2006) and treat a setting in which aggregate consumption is stationary (i.e., a time-invariant function of dividends) but individual endowments are not.¹⁴

We proceed as follows. We first present an abstract version of the economic structure of the laboratory environment (endowment processes, preferences, risk, etc.) and discuss the theoretical predictions. We then explain how the abstract economic structure was obtained in the laboratory (with elements such as inducing of preferences, ensuring stationarity of the risk in the economy, organization of trade, markets, and communication).

A. *The Structure of The Laboratory Economy*

We consider an infinite-horizon economy with a single perishable consumption good in each time period. In the experiment, the consumption good is cash so we use the terms “consumption” and “cash” interchangeably here. How we make cash “perishable” and how we make the laboratory economy infinitely lived will be discussed below. In each period there are two possible states of nature, H (high) and L (low), which occur with probabilities π and $1 - \pi$ independent of time and past history. Two long-lived assets are available for trade: (i) a *Tree* that pays a stochastic dividend d_T^H when the state is H and d_T^L when the state is L , and (ii) a (*consol*) *Bond* that pays a constant dividend $d_B^H = d_B^L = d_B$ each period.¹⁵ We assume $d_T^H > d_T^L \geq 0$ and normalize so that the Bond and Tree have the same expected dividend: $d_B = \pi d_T^H + (1 - \pi)d_T^L$. Note that the dividend processes are stationary

in levels. (In the experiment proper, we choose $\pi = 1/2$, $d_T^H = 1$, $d_T^L = 0$, and $d_B = 0.50$, with all payoffs in dollars.)

There are n agents, where n is even (in the experiments n will be between 12 and 30). Each agent i has an initial endowment b_i of Bonds and τ_i of Trees, and also receives an additional private flow of income $e_{i,t}$ (possibly random) in each period t . Write $b = \sum b_i$, $\tau = \sum \tau_i$, and $e = \sum e_i$ for the social (aggregate) endowments of Bonds, Trees, and additional income flow. We assume that the social income flow e is stationary – that is, a time-invariant function of dividends (in the experiment proper it will be constant) – so that aggregate consumption $bd_B^\sigma + \tau d_T^\sigma + e$ is also stationary (σ indexes the state), but we impose *no* restriction on individual endowments. (As noted earlier, we wish to ensure that in the experimental setting subjects have a reason to trade each period.)

We induce the following preferences. Agent i maximizes discounted expected lifetime utility for infinite (stochastic) consumption streams

$$U_i(\{c_t\}) = E \left[\sum_{t=1}^{\infty} \beta^{t-1} u_i(c_t) \right],$$

where c_t is (stochastic) consumption at time t . We assume that the period utility functions u_i are smooth, strictly increasing, strictly concave, and have infinite derivative at zero (so that optimal consumption choices are interior), but we make no assumptions as to functional forms. Note that agent endowments and utility functions may be heterogeneous but that all agents use the same constant discount factor β , which we induce to equal $5/6$.

In each period t agents receive dividends from the Bonds and Trees they hold (as well as potential income flow), trade their holdings of Bonds and Trees at current prices, use the proceeds together with their endowments to buy a new portfolio of Bonds and Trees, and consume the remaining cash. How exactly agents buy and sell in our laboratory economy will be explained below. Here, we follow the theory and assume that agents take as given the current prices of the Bond $p_{B,t}$ and the Tree $p_{T,t}$ (both of which depend on the current state), make forecasts of (stochastic) future asset prices $p_{B,t'}$ and $p_{T,t'}$ for each $t' > t$, and optimize subject to their current budget constraint and their forecast of future asset prices. (More directly: agents optimize subject to their forecast of future consumption conditional on current portfolio choices.) At a *Radner equilibrium* (Radner (1972)), markets for consumption and assets clear at every date and state and all price forecasts are correct (“perfect foresight”). This is not quite enough for equilibrium to be well defined because it does not rule out the possibility that agents acquire more and more debt, delaying repayment further and further into the future – and never in fact repaying it. In order for equilibrium be well defined, such schemes must be ruled out. Levine and Zame (1996, Magil and Quinzii (1994), and Hernandez and Santos (1996) show that this can be done in a number of ways. Levine and Zame (1996) show that all “reasonable” ways lead to the same equilibria; the simplest is to require that debt not become unbounded.¹⁶ (In the experimental setting, we forbid short sales so debt is necessarily bounded.)

B. Predictions

Our predictions derive from an analysis of the equilibrium in this economy. We assume that this is a Radner (perfect foresight) equilibrium and, as is universal in the literature, that this Radner equilibrium exists and (because markets are (potentially) dynamically complete) coincides with Walrasian equilibrium, so that equilibrium allocations are Pareto optimal. These assumptions are not innocuous, but, as noted before, the familiar version of the Lucas model *begins* with the assumption of a representative agent equilibrium, and the existence of a representative agent assumes Pareto optimality. Thus, from a theoretical point of view, all that we are assuming is *subsumed* in the familiar version. Whether a Radner equilibrium obtains that generates Pareto optimal allocations is ultimately an empirical issue. Verification that a particular equilibrium obtains with associated welfare properties has occupied economists who work on market experiments ever since the seminal paper of Smith (1965).

Despite allowing for heterogeneity and without making any assumptions about functional forms for preferences, the theory makes testable *quantitative* predictions about individual consumption (which, in the experiments, will be end-of-period cash holdings), prices, and trading patterns. Some of these predictions take a particularly simple form when the specific parameters are as in our experiments. The pricing predictions should be entirely familiar in the context of the usual Lucas model with a representative agent having constant relative

risk aversion (CRRA) utility (though our predictions do not rely on CRRA). We go beyond pricing because we formulate predictions about allocations as well.

PREDICTION 1: *There is trade in each period.*

For investors in our economy to reach Pareto optimal allocations, they need to smooth consumption to the point that their consumption is stationary and perfectly correlated with aggregate consumption. Since their income fluctuates between periods (in fact, we let it fluctuate between odd- and even-numbered periods), they need to un-do these fluctuations, and they can do so (only) through trade in the financial markets. This means that trading volume must always be positive. Volume may be larger in the first period (when investors trade to their long-run average optimal securities holdings), but beyond the first period they trade to offset income fluctuations.

PREDICTION 2: *The Tree is always cheaper than the Bond.*

Because markets are dynamically complete, agents can trade to Pareto optimal allocations. Pricing in a Pareto-optimal allocation can be derived using the representative agent approach. Indeed, a representative agent exists, and prices should be such that the representative agent is willing to hold the supply of assets and consume no more or less than the aggregate dividend.

In an economy with heterogeneous agents, the preferences of the representative agent are hard to derive. These preferences may not even look like those of any individual agent.

However, since each period we only have two possible states (the state σ is either high H or low L), things simplify dramatically.

To see this, fix an individual agent i ; write $\{c_{i,t}\}$ for i 's equilibrium consumption stream. Because of Pareto optimality and of stationarity of aggregate dividends, consumption depends only on the state and not on time (see Appendix A for details), so we suppress the time index t . Write i 's first-order condition for optimality

$$p_{A,t}^\sigma = \beta \left\{ \pi \left[\frac{u'_i(c_i^H)}{u'_i(c_i^\sigma)} \right] (d_A^H + p_{A,t+1}^H) + (1 - \pi) \left[\frac{u'_i(c_i^L)}{u'_i(c_i^\sigma)} \right] (d_A^L + p_{A,t+1}^L) \right\},$$

where superscripts index states and subscripts index assets, time, and agents in the obvious way. We can write this in more compact form as

$$p_{A,t}^\sigma = \beta E \left\{ \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} \right] (d_A + p_{A,t+1}) \right\}, \quad (1)$$

for $\sigma = H, L$ and $A = B, T$. Equality of the ratios of marginal utilities across *all* agents, which is a consequence of Pareto optimality, implies that (1) is independent of the choice of agent i , and hence we could write (1) in terms of the utility function of a representative agent.

Let $\mu = u'_i(c_i^L)/u'_i(c_i^H)$ be the marginal rate of substitution in the low state for consumption in the high state. Pareto optimality guarantees that μ is independent of which

agent i we use, so it is the marginal rate of substitution of any individual *and* that of the representative agent. (Note that risk aversion implies $\mu > 1$.) Consequently,

$$\begin{aligned} p_A^H &= \beta \left[\pi(d_A^H + p_A^H) + (1 - \pi)(d_A^L + p_A^L)\mu \right] \\ p_A^L &= \beta \left[\pi(d_A^H + p_A^H)(1/\mu) + (1 - \pi)(d_A^L + p_A^L) \right]. \end{aligned}$$

Solving yields

$$\begin{aligned} p_A^H &= \left(\frac{\beta}{1 - \beta} \right) \left[\pi d_A^H + (1 - \pi)d_A^L \mu \right] \\ p_A^L &= \left(\frac{\beta}{1 - \beta} \right) \left[\pi d_A^H (1/\mu) + (1 - \pi)d_A^L \right]. \end{aligned} \quad (2)$$

Specializing to the parameters of the experiment $d_T^H = 1, d_T^L = 0, d_B^H = d_B^L = 0.5$, and $\beta = 5/6$ yields

$$\begin{aligned} p_B^H &= (2.5)(1 + \mu)/2 \\ p_B^L &= (2.5)(1 + \mu)/2\mu \\ p_T^H &= 2.5 \\ p_T^L &= 2.5/\mu. \end{aligned} \quad (3)$$

The important thing to note here is that

$$p_B^\sigma > p_T^\sigma$$

in each state σ . That is, the Bond is always priced above the Tree. Intuitively, this is because the “consumption beta” of the Tree is higher and hence is discounted more (relative to expected future dividends). The consumption beta of a security is the covariance of its future dividends with aggregate future consumption. Bond dividends are deterministic, while those of the Tree increase with aggregate consumption. Hence, the consumption beta of the Tree is higher than that of the Bond.

Notice also that, under our parametrization, $p_T^H = 2.5$, that is, the price of the tree in the high state is independent of risk attitudes. In addition, $p_B^H/p_B^L = p_T^H/p_T^L$, that is, the ratios of asset prices in the two states are the same.

PREDICTION 3: *Asset prices are perfectly correlated with fundamentals.*

From (3) it follows that

$$p_A^H/p_A^L = \mu. \tag{4}$$

Consequently, prices are perfectly correlated with the state – they are higher in the good state (H) and lower in the bad state (L). Significantly, prices *only* change if the state changes: if the state is the same in two consecutive periods, prices do not change.

Why prices are higher in the good state can be understood most clearly by thinking about the representative agent. In state H , aggregate consumption supply is high, so high prices (low returns) must be in place to temper the representative agent's desire to save (buy). The opposite is true for state L : aggregate consumption is low, so low prices (high returns) temper the representative agent's desire to borrow (sell).

PREDICTION 4: *The more the Tree trades at a discount relative to the Bond, the greater is the difference in the prices of both securities across states.*

More precisely, cross-sectional and intertemporal features of asset prices reinforce each other.¹⁷ The discount of the Tree relative to the Bond increases because risk aversion rises. As a result, the Tree and Bond prices move more extremely with fundamentals. Mathematically,

$$\text{cov}(p_B^\sigma - p_T^\sigma, p_A^H - p_A^L) > 0,$$

for $\sigma = H, L$ and $A = B, T$, with covariance computed based on sampling across cohorts of agents (economies), keeping everything else constant. "Everything else" means: initial endowments, private income flows, asset structure, outcome probabilities, as well as impatience β . Economies are therefore distinguishable at the price level only in terms of the risk aversion (embedded in μ) of the representative agent.¹⁸

PREDICTION 5: *Expected returns vary across states.*

From prices and dividends of the Tree and Bond as well as the state transition probabilities (from a low state one moves to a high state or remains in the low state with equal probability), one can readily compute the expected returns on the Tree and the Bond. Simple algebraic manipulation then allows one to express the difference across high and low states of the expected return on the Tree ($E[R_T|H] - E[R_T|L]$) and Bond ($E[R_B|H] - E[R_B|L]$) as follows:

$$\begin{aligned} E[R_T|H] - E[R_T|L] &= \pi(1 - \mu) + (1 - \pi)\left(\frac{1}{\mu} - 1\right) + d_B(1 - \mu)\frac{1}{p_T^H}, \\ E[R_B|H] - E[R_B|L] &= \pi(1 - \mu) + (1 - \pi)\left(\frac{1}{\mu} - 1\right) + d_B(1 - \mu)\frac{1}{p_B^H}. \end{aligned}$$

Because $\mu > 1$, all terms in both expressions are negative, and hence the expected return on both assets is higher in “busts” (when the state is low) than in “booms” (when the state is high). Lucas (1978) emphasized return predictability as one of the natural equilibrium features of the model and contrasted this with the martingale version of EMH.¹⁹

PREDICTION 6: *The equity premium is countercyclical.*

The difference between the expected return on the risky security (the Tree) and the expected return on the (relatively) risk-free security (the Bond) is the *equity premium* (Mehra and Prescott (1985)).²⁰ Because asset prices and dividends are stationary, equity premia are stationary as well and depend only the state. Simple computations then allow us to explicitly derive the equity premia in the high and low states. These computations rely on two facts:

(i) the expected dividends are the same for both assets and equal to d_B , and (ii) for each asset A , $p_A^H = \mu p_A^L$. We obtain

$$\begin{aligned}
 E^H &= \frac{\pi p_T^H + (1 - \pi)p_T^L + d_B - p_T^H}{p_T^H} - \frac{\pi p_B^H + (1 - \pi)p_B^L + d_B - p_B^H}{p_B^H} \\
 &= d_B \left(\frac{1}{p_T^H} - \frac{1}{p_B^H} \right) \\
 E^L &= \frac{\pi p_T^H + (1 - \pi)p_T^L + d_B - p_T^L}{p_T^L} - \frac{\pi p_B^H + (1 - \pi)p_B^L + d_B - p_B^L}{p_B^L} \\
 &= d_B \left(\frac{1}{p_T^L} - \frac{1}{p_B^L} \right) = d_B \left(\frac{\mu}{p_T^H} - \frac{\mu}{p_B^H} \right) = \mu E^H.
 \end{aligned}$$

Note that both equity premia are positive. The difference across states is

$$E^H - E^L = (1 - \mu)E^H.$$

This difference is strictly negative (because $\mu > 1$), so the equity premium is countercyclical: it is lower in the high state than in the low state. Note that countercyclicity provides the right incentives: when dividends are low the equity premium is high, so investors buy risky trees rather than consuming scarce dividends; when dividends are high the equity premium is low, so investors prefer to consume rather than engage in risky investment.

PREDICTION 7: *Agents smooth consumption over time.*

This is an immediate consequence of Pareto optimality, the condition that guarantees the existence of a representative agent. Despite the fact that individual incomes fluctuate over time, in equilibrium agents smooth out this fluctuation by trading in securities. Income

fluctuations do not impact available aggregate consumption, and hence agents should be able to trade them away.

PREDICTION 8: *Agents hedge fundamental risk.*

Equilibrium trading strategies involve (net) selling assets when income is low and (net) re-purchasing assets when income is high. Equilibrium prices change with fundamentals (aggregate dividends), and since fundamentals are random, prices are random. Because prices are random, agents do not know at what prices they will be able to carry out these repurchases, that is, there is price risk. However, they have the opportunity to hedge this price risk. Indeed, the dividend of the Tree is (perfectly) correlated with aggregate dividends, and hence with prices. By buying Trees, beginning-of-period cash balances are high when the dividend of the Tree is high. When this happens, aggregate dividends, and hence prices, are high too. The increase in cash balances allows agents to repurchase assets despite the increase in prices. The opposite obtains when the dividend of the Tree is low: cash balances at the beginning of the period are (relatively) low, but that is no concern because repurchase prices will be low as well. As such, Trees allow agents to hedge fundamental risk, and hence price risk. Consequently, but perhaps counterintuitively, agents should buy Trees in periods when private income is low and should sell when private income is high. In fact, it can be shown that hedging increases with risk aversion. This is because price risk, measured as the difference in prices across H and L states, increases with risk aversion (reflected in μ).

We reiterate that, in the Lucas equilibrium, prices move only because of aggregate dividends. So, fundamental risk is the sole source of price risk. Any residual price movement is incompatible with the Lucas model. Residual price risk may emerge in equilibrium because agents do not know how prices relate to the fundamental state of the economy, that is, when agents have imperfect foresight. We elaborate on this point below. From the point of view of the agents, residual price risk and the *residual price forecasting risk* we alluded to before are the same. Residual price forecasting risk is the residual randomness in prices given a fundamental state (i.e., a level of aggregate dividends).

Hedging is usually associated with Merton's intertemporal asset pricing model (Merton (1973a)) and is the core of modern derivatives analysis (Black and Scholes (1973), Merton (1973b)), but it also forms an integral part of the trading predictions of the Lucas model.

In summary, our implementation of the Lucas model predicts that securities prices differ cross-sectionally depending on consumption betas (the Tree has the higher beta), while intertemporally, securities prices move with fundamentals (dividends of the Tree). The two predictions reinforce each other: the bigger the difference in prices across securities, the larger the intertemporal movements. These predictions are the same as in the original Lucas model. But in our setting, there is agent heterogeneity, and we can say more. Specifically, investment choices should be such that consumption (cash holdings at the end of a period) across states becomes perfectly rank-correlated between agents (or even perfectly correlated, if agents have the same preferences). Likewise, consumption should be smoothed across periods with and

without income. Investment choices are sophisticated: they require, among other things, that agents hedge price risk by buying Trees when experiencing an income shortfall (and selling Bonds to cover the shortfall), and selling Trees in periods of high income (while buying back Bonds).

For illustration, Appendix B provides explicit solutions for equilibrium prices, holdings, and consumption when taking the parameters as in the experiment and assuming that all agents display identical CRRA, which we vary from 0.2 to 1.0. There is nothing special about these particular choices of risk aversion; we offer them solely for comparison purposes. We note that risk aversion in the range 0.2 to 0.5 is consistent with the experimental findings of Holt and Laury (2002) and Bossaerts and Zame (2008).

IV. Implementation in The Laboratory

As we have already noted, implementing the Lucas economy in the laboratory involves three difficulties:

- (a) The Lucas model has an infinite horizon and assumes that agents discount the future.
- (b) The Lucas model assumes that agents prefer to smooth consumption over time.
- (c) The Lucas economy is stationary.

In our experiment, we use the standard solution (Camerer and Weigelt (1996)) to resolve issue (a), which is to randomly determine if a period is terminal. This ending procedure

induces discounting with a discount factor equal to the probability of continuation. We set the termination probability equal to $1/6$ so the continuation probability (and induced discount factor) is $\beta = 5/6$. In mechanical terms: after the markets in period t close we roll a 12-sided die; if it comes up either 7 or 8, we terminate, otherwise we move on to a new period.

To resolve issue (b), we make end-of-period individual cash holdings *disappear* in every period that is not terminal; only securities holdings carry over to the next period. If a period *is* terminal, however, securities holdings perish and cash holdings are credited; participants' earnings are then determined entirely by the cash they hold at the end of this terminal period. To see that this has the desired implication for preferences, note that the probability that a given replication terminates in period t is the product of $(1 - \beta)$ (the probability that it terminates in period t , conditional on not having terminated in the first $t - 1$ periods) times β^{t-1} (the probability that it does not terminate in the first $t - 1$ periods). Hence, assuming expected utility, each agent maximizes

$$\sum_{t=1}^{\infty} (1 - \beta) \beta^{t-1} E[u(c_t)] = (1 - \beta) E \left[\sum_{t=1}^{\infty} \beta^{t-1} u(c_t) \right].$$

Of course, the factor $(1 - \beta)$ has no effect on preferences.²¹

So, while we will not literally have consumption every period, agents in our setting face the same optimization problem as if they actually had to consume every period and wanted to smooth consumption across periods. Thus, our experiment is designed so that, for subjects

to optimize, they need to solve the same problem in the laboratory as in the theory. When choices for the two optimization problems are the same, the Lucas model predicts that prices will also be the same, because in the Lucas model, prices *derive* from choices, by means of the stochastic Euler equations.

It is less obvious how to resolve problem (c). The problem is not with the dividends and personal income but with the termination probability. In principle, simply announcing a constant termination probability should do the trick – because each period is equally likely to be terminal. However, if the probability of termination is in fact constant (and independent of the current duration), then the experiment could continue for an arbitrarily long time. In particular, there would be a nonnegligible probability that the experiment would continue much longer than a typical session. It is clear that subjects understand this: in our own pilot experiments, subject beliefs about the termination probability increased substantially as the end of the session approached.

To deal with this problem we employ a simple termination rule: we announce that the experimental session will last until a pre-specified time and there will be as many replications of the (Lucas) economy as possible within this time frame. If a replication ends at least 10 minutes before the announced ending time of the session, a new replication begins; otherwise, the experimental session ends. If a replication is still running 10 minutes before the announced ending time of the session, we announce before trade opens that the current period will be either the last one (if our die turns up 7 or 8) or the next-to-last one (for all other values of the die). In the latter case, the next period is the terminal period, with

certainty, so subjects would keep the cash they received through dividends and income for that period. Note that there should be no trade in the terminal period because assets perish at the end and hence are worthless. We nevertheless opened markets, to verify whether subjects understood this. Invariably, they did. In Appendix C, we reproduce the timeline that we provided alongside the instructions to facilitate comprehension.

To see that equilibrium prices remain the same regardless of whether the new termination protocol is applied or whether termination is perpetually determined with the roll of a die, consider an agent's optimization problem in period t , which is terminal with probability $1 - \beta$ and penultimate with probability β : maximize $(1 - \beta)u(c_t^\sigma) + \beta E[u(c_{t+1})]$ subject to the standard budget constraint. The first-order conditions for asset A are

$$(1 - \beta)u'(c_t^\sigma)p_{A,t}^\sigma = \beta E[u'(c_{t+1})d_{A,t+1}].$$

The left-hand side is the expected marginal utility from keeping cash worth one unit of the security, while the right-hand side is the expected marginal utility from buying the unit; optimality implies equality. Rearranging yields

$$p_{A,t}^\sigma = \left(\frac{\beta}{1 - \beta} \right) E \left[\frac{u'(c_{t+1})}{u'(c_t^\sigma)} d_{A,t+1} \right].$$

Because aggregate consumption and dividends are i.i.d., this can be written as an infinite series:

$$p_{A,t}^{\sigma} = \sum_{\tau=0}^{\infty} \beta^{\tau+1} E \left[\frac{u'(c_{t+\tau+1})}{u'(c_t^{\sigma})} d_{A,t+\tau+1} \right].$$

The series telescopes, so can be written recursively as in (1), which is the equation that applies to the infinite-horizon case. Hence, our two-period termination protocol produces pricing *as if* the economy were to continue forever.²²

In short, to utilize the two-hour session time optimally, we run as many replications as possible. Whenever a replication terminates, if there is sufficient time left in the session we initiate a new replication. Thus, the termination protocol is only applied to at most one of the replications in a session. If a replication ends naturally (i.e., through the roll of the die) close to the 10-minute mark before the end of a session, we do not start a new replication.

Overall, our design may look intimidating, but the task we expected the subject to perform is actually quite intuitive. Each period, the subject enters with holdings of the Tree and Bond from a prior period, and receives cash, in the form of income (potentially) and dividends. The subject then has to decide how to redistribute her wealth across cash and assets. She knows that there is a chance (1/6) that the replication ends that period, at which point she earns the cash she is holding but the assets she still has in her portfolio become worthless. With complementary chance (5/6), we move to a subsequent period and the subject forfeits her cash while rolling over assets. The latter generate new cash (in the form of dividends) in the new period. So, the task we ask the subject to perform is to optimally

balance holdings of cash and assets. Assets are risky, but by holding Bonds, dividend risk can be eliminated. So, some combination of Trees and Bonds seems to be optimal. To ensure subjects understand the task, the instructions (reproduced in Appendix C) contain a table that specifies, for two different strategies, the effect on cash and in turn on potential earnings.

There is one further difficulty that we have not discussed: default. In the (finite- or infinite-horizon) Radner model, assets are simply promises – selling an asset, that is, borrowing, entails a promise to repay in the future. However, in the model, nothing enforces these promises – that they are kept is simply part of the definition of equilibrium. If nothing enforces these promises in the laboratory, then participants could (and in our experience, would) simply make promises that they could not keep. One way to address this problem is to impose penalties for default, that is, for failing to keep promises. In some sense that is what Radner equilibrium implicitly presumes: there are penalties for default and these penalties are so great that no one ever defaults.

However, imposing penalties in the laboratory is highly problematic. First, what should the penalties be? The rules governing experimentation with human subjects prevent us from forcing subjects to pay from their own pockets, and excluding subjects from further participation in the experiment would raise a host of problems following such exclusion – to say nothing of the fact that neither of these penalties might be enough to guarantee that default would not occur and to make it common knowledge that default would not occur. Moreover, the above speaks only to intentional default, but what about *unintentional default*,

that is, mistakes? And what about plans that would have led to default in circumstances that might have occurred but did not? Further, what about the fact that the mechanisms for discouraging default might change behavior in other, unexpected, ways?

There is no simple solution to this problem because it is not a problem confined to the laboratory. Radner equilibrium effectively prohibits default but it is entirely silent about how this prohibition is to be enforced. As Kehoe and Levine (1993) and Geanakoplos and Zame (2014) (and others) point out, mechanisms for dealing with default may eliminate default, but only at the cost of other distortions.

Our solution in the laboratory is simply to prohibit short sales (negative holdings) of assets. This creates a potential problem because our theoretical analysis presumes that it is always possible for any agent to buy or sell an infinitesimal additional quantity of either asset, but if an agent's current holding of an asset were zero he could not sell it and if his consumption and portfolio were both zero he could not buy it. However, so long as agents do not bump up against the zero bound, our analysis remains correct. In the actual experimental data, the number of agents who bumped up against the zero bound was quite small.

The need to bar short sales explains why we use an instantiation of the Lucas economy in which the Bond is in positive net supply: risk-tolerant subjects could merely reduce their holdings of Bonds rather than having to sell short (which is not permitted). Because both assets are in positive supply, our economy is, strictly speaking, a Lucas *orchard* economy

(Martin (2011)), but the predictions discussed above are not different in a model in which the Bond is in zero net supply.

Because income and dividends vary over time and across states, and because cash disappears at the end of each nonterminal period, those who are optimizing or nearly optimizing must trade often. As we shall see, trading volume was indeed substantial.²³ Trading took place through an anonymous, electronic continuous open-book system. The trading screen, part of software called Flex-E-Markets, was intuitive, requiring little instruction.²⁴ Subjects quickly familiarized themselves with key aspects of trading in the open-book mechanism (bids, asks, cancellations, transaction determination protocol, etc.) through one mock replication of our economy during the instructional phase of the experiment. A snapshot of the trading screen is reproduced in Figure 1. All accounting and trading was done in U.S. dollars. Thus, subjects did not have to convert from imaginary experiment money to real-life currency.

Figure 1 HERE

As explained above, within each experimental session, we conduct as many replications as possible within the time allotted. To minimize wealth effects, we pay for only a fixed number (2) of the replications, randomly chosen after conclusion of the experiment.²⁵ However, we do not view wealth effects as important in this context in any event, since there is no assumption that subject preferences are constant across replications within a given session – and they are

certainly not constant across sessions, since the populations of subjects in different sessions were disjoint.

Finally, we allocate assets and income processes as follows:

- There are an even number $n = 2m$ of agents; agents $i = 1, \dots, m$ are of Type I, while agents $i = m + 1, \dots, 2m$ are of Type II.
- Type I agents are endowed with asset holdings $b_I = 0, \tau_I = 10$ and have income $e_{I,t} = 15$ when t is even and $e_{I,t} = 0$ when t is odd.
- Type II agents are endowed with asset holdings $b_{II} = 10, \tau_{II} = 0$ and have income $e_{II,t} = 15$ when t is odd and $e_{II,t} = 0$ when t is even.

V. Results

Table I provides specifics of the six experimental sessions, each of which contained several replications. The number of participants ranged from 12 to 30. Three sessions were conducted at Caltech, two at UCLA, and one at the University of Utah, thus providing us with evidence from a variety of subject pools (subjects were undergraduate and graduate students as well as staff from these institutions; they were not selected based on their background in economics or finance). In all there were 15 replications, totaling 80 periods. Whenever the end of the experiment occurred during a replication (starred sessions), our novel termination protocol was applied: in the terminal period of these replications, participants knew

for certain that it was the last period. In the other (unstarred) sessions, the last replication occurred sufficiently close to the end of the experiment that a new replication was not begun, so our termination protocol was not applied.

Table I HERE

We now discuss the extent to which each prediction is borne out in the data. We first restate the prediction (in the same order as before) and then evaluate the evidence.

PREDICTION 1: *There is trade in each period.*

Table II reports average trading volume per period (excluding terminal periods during which there was no trade). Trading volume in periods 1 and 2 is significantly higher, reflecting the need for agents to move from initial holdings to their steady-state holdings. In the theory, subsequent trade takes place only to smooth consumption across odd and even periods.²⁶ Volume in the Bond is significantly lower in periods 1 and 2. This appears to be an artifact of the few replications when the state in period 1 was low, which deprived Type I participants (who were endowed with 10 Trees and had no personal income in odd periods) of cash. In principle, Type I participants should have been able to sell enough Trees to buy Bonds, but it appears that they did not manage to complete all the necessary trades in the allotted time (four minutes). On average, 23 Trees and 17 Bonds were traded per period. Since the average supply was 210 Trees and 210 Bonds and the average period was three to four

minutes long, this means that roughly 10% of available securities were traded each period and that one transaction occurred roughly every five seconds.

Table II HERE

PREDICTION 2: *The tree is always cheaper than the bond.*

Table III displays average period transaction prices. Table IV presents the results by the period's state (High if the dividend of the Tree was \$1, Low if it was \$0). Consistent with the theory, the Bond is priced above the Tree, with the price differential averaging about \$0.50. For the moment, we refrain from making statements about the statistical significance. We do so later when we discuss a proper modeling of the dynamics of (average-period) prices. Indeed, the time series exhibit peculiar properties, as we discuss next.

Table III HERE

PREDICTION 3. *Asset prices are perfectly correlated with fundamentals.*

Prices move with fundamentals, but the correlation is *far from* perfect. Figure 2 plots the evolution of (average) prices over time, arranged chronologically by experimental sessions (numbered as in Table I); replications within a session are concatenated.²⁷ The plot reveals that prices are volatile. In theory, prices should move only because of variability in economic fundamentals, which in this case amounts to changes in the dividend of the Tree. Prices should be high in high states, and low in low states, and these levels should be determined

only by the state, not by, for example, passage of time or whether the period is odd or even. In reality, a large fraction of price movements is unrelated to fundamentals; following LeRoy and Porter (1981) and Shiller (1981), we refer to this as *excessive volatility*.

Figure 2 HERE

Some price drift can be detected, but formal tests reported below reveal that the drift is entirely due to the impact of states on prices, and the particular sampling of states across the sessions. That is, after accounting for changes in states, price changes are white noise.

According to the theory, p_T^H should equal 2.50 irrespective of risk aversion (of the representative agent) and all other prices (across states, securities) derive from this number by means of a single parameter. As Table IV shows, p_T^H is higher than predicted: it is predicted to be 2.50 but is actually 2.91. Similarly, p_T^H/p_T^L is greater than p_B^H/p_B^L while the two ratios should be equal. As before, we refrain from taking a position on the statistical significance of these deviations. (Period-average) transaction prices are not i.i.d., so we cannot rely on standard t tests to determine significance.

Table IV HERE

PREDICTION 4: *The more the Tree trades at a discount relative to the Bond, the greater is the difference in the prices of both securities across states.*

The prediction is that the differential in prices between high and low states should correlate positively with the difference between the Bond price and the Tree price, that is, the discount

of the Tree price relative to the Bond price. The correlation is to be taken across economies, that is, across replications. Table V displays correlations of the average discount on the Tree price relative to the Bond price (regardless of state) and the average difference between prices of the Tree or of the Bond across states. Each observation corresponds to one replication, so there are 15 observations in total. Consistent with the theoretical prediction, the correlations are positive, though the estimate is insignificant for the Bond.

Table V HERE

Formal Statistics for Predictions 2-3-4

At this point, we pause and discuss statistical significance. To assess significance of price differences across states, we regress period transaction price levels on the state (=1 if high; 0 if low). To adjust for time-series dependence that one might perceive in Figure 2, we originally added session dummies and a time trend (period number). In addition, to gauge the effect of our session termination protocol, we added a dummy for periods when we announce that the session is about to come to a close, and hence the period is going to be either the penultimate or last one, depending on the roll of the die. Lastly, we originally also added a dummy for even periods.

However, closer inspection of the properties of the error term revealed substantial dependence over time, as evident from a Durbin-Watson test.²⁸ Further model specification analysis was performed, to ensure that the error term became properly behaved. This re-

vealed that the best model required first differencing of price changes. The highest adjusted R^2 was obtained for a model that predicted price changes across periods as a result of only the change in the state. See Table VI.²⁹ The regression does not include an intercept or period dummies: average price changes are insignificantly different from zero and independent of the period once the change in state is accounted for. This implies, among other things, that the apparent drift in the display of the price data (Figure 2) is a visual illusion (besides being only intra-session, not within-replication), driven entirely by the actual sampling of the state. The autocorrelations of the error terms are now acceptable (comfortably within two standard errors from zero).

For the Tree, the effect of a change in state from low to high is substantial (\$0.19) and significant ($p < 0.05$). The effect of a change in state on the Bond price is smaller (\$0.10), and insignificant ($p > 0.05$). Both confirm the theoretical prediction that prices should be determined by the state.

Table VI HERE

The excess volatility of prices is plainly evident from Table VI. Fundamentals (changes in the state) explain only 18% of the variability of the Tree prices ($R^2 = 0.18$); 82% of the price variance is unexplained.³⁰ The situation is even worse for the Bond: 96% of the variance of Bond price changes is unexplained by changes in the state. The Lucas model instead predicts that fundamentals should explain *all* of the variability of asset prices. It is worth highlighting that the unexplained variability is essentially noise; in particular, it is unrelated

to the subject cohort, because session dummies were insignificant. Overall, the regressions in first differences show that, consistent with the Lucas model, fundamental economic forces are behind price changes – significantly so for the Tree. But at the same time, prices are excessively volatile, with no discernible drift.

Prediction 5: *Expected returns vary across states.*

The theory states that the expected returns on the two securities should vary depending on the state. Specifically, expected returns are higher in the low than in the high state, which is to say that expected returns are countercyclical. In Table VII we find that the results are consistent with the theory: the difference in average returns across high and low states is negative for both securities, though insignificant (at 5% level) for the Bond. The countercyclicity of expected returns implies that returns *are* predictable, in violation of traditional accounts of EMH.

Table VII HERE

PREDICTION 6: *The equity premium is countercyclical.*

Table VII demonstrates that the equity premium is countercyclical. But the theory also requires the equity premium to be positive, while it is not in “boom” periods: the average return on the Tree is below that of the Bond when the state is high, and the difference is highly significant. In the low state, and overall, the equity premium is positive but not significant. This return anomaly in the high state is in sharp contrast with the general

correctness of price levels. This paradoxical situation may be understood in view of the noise in prices, and since returns are ratios of prices (and dividends), the noise is amplified in the returns. As a result, the fit of the Lucas model is much worse when one considers returns rather than price levels.

PREDICTION 7: *Agents smooth consumption over time.*

In the Lucas equilibrium, consumption choices are Pareto optimal. This means, in particular, that agents of both types should trade to holdings that generate high consumption in high states, and low consumption in low states. Table VIII displays the average amount of cash (consumption) per type in high vs. low states.³¹ Consistent with the theoretical prediction, consumption is positively correlated across types.

To gauge the significance of this finding, Table VIII also displays, in parentheses, the consumption (cash) levels that agents could have reached if they were not allowed to trade. These are the consumption levels under autarky. Note that autarky consumption levels are *anticorrelated*. Through trading, the average Type I and Type II agents managed to move correlation of their consumptions from negative to positive. This suggests economically significant Pareto improvements, consistent with the Lucas model.

Table VIII HERE

Pareto optimality also implies that subjects should be able to perfectly offset income differences across odd and even periods. Table VIII demonstrates that our subjects did

indeed manage to smooth consumption substantially; the outcomes are far more balanced than under autarky (numbers in parentheses; averaged across high and low states, excluding periods 1 and 2), again implying substantial Pareto improvements.³²

If one were willing to entertain the assumption that utilities of our subjects are homothetic, Pareto efficiency would suggest a stronger prediction than positive correlation of consumption across states, or smoothing of income across odd and even periods. Under homothetic utilities, consumption shares should be independent across states and across periods. Table VIII displays the results of a formal test of equality of the consumption share of the average Type I subject across states and periods. The share of total consumption (total cash available) that the average Type I subject chose at the end of each period was computed and a two-way analysis of variance (ANOVA) was applied, with state (high/low) and period (odd/even) as potential factors determining variability in this consumption share, allowing for interaction between state and period. A mixed-effects approach was used, to accommodate differences in consumption shares across replications due to differences in drawing of the state in the first period and in subject cohort.

Table VIII shows that neither the state nor the nature of the period (nor their interaction) are significant ($p > 0.05$) factors in explaining the variability of the consumption share of the average Type I subject across periods. As such, the apparent violations of the prediction of equal consumption shares across states/periods implied by the average consumption levels reported in Table VIII are due solely to sampling error.

The finding is rather striking, because the assumption of homothetic preferences is questionable. Yet, our empirical results suggest that the assumption can be maintained as far as the choices of the average subject of Type I (and by implication, of Type II) are concerned.

PREDICTION 8: *Agents hedge fundamental risk.*

The above results show that our subjects (on average) managed to move towards the Pareto-optimal equilibrium consumption patterns of the Lucas model. However, they did not resort to hedging as a means to ensure those patterns.

Table IX lists average asset holdings across periods for Type I subjects (who received income in even periods). They were net sellers of assets in periods of income shortfall (see the “Total” row). But unlike in the theory, they *decreased* Tree holdings in low-income periods and *increased* them in high-income periods. Only in period 9 is there some evidence of hedging: Type I subjects on average *bought* Trees while they were income-poor (period 9’s holding of Trees is higher than period 8’s).

This finding should put our trading volume numbers in perspective. As noted earlier, trade volume was substantial (see Table II). However, the volume of trade is much lower than predicted by the Lucas model. In Appendix B, we show that even for (constant) relative risk aversion as low as 0.2, average turnover should be about 20% of average holdings for the Tree, and over 50% for the Bond. Keep in mind, however, that these predictions presume that subjects follow the sophisticated hedging strategy of Prediction 8, buying Trees when in need of cash. Without this strategy, the volume needed to smooth consumption is substantially

lower. Fewer assets have to be traded since Tree purchases (required for hedging) need not be offset by extra Bond sales (required to smooth consumption).

If their actions truly reflect their beliefs, then the absence of hedging suggests that subjects did not expect prices to co-move with the dividend on the Tree. Indeed, if prices move at random relative to fundamentals, then price risk cannot be hedged by buying Trees, unlike in the Lucas equilibrium. Income shortfalls can be covered as well by selling Trees as by selling Bonds. This observation turns out to form a crucial basis for a conjecture as to why prices exhibit excess volatility; we elaborate in Section VII.

Table IX HERE

Subject-Level Differences

There are substantial individual differences in portfolio choices. Table X illustrates how three Type I subjects end up holding almost opposing portfolios of Trees and Bonds. Subject 7 increased his holdings of Trees over time. Significantly, this subject bought Trees in periods with income shortfall (odd periods), effectively implementing the hedging strategy of the theory. Subject 5 is almost a mirror image of subject 7, though s/he did not resort to hedging. Subject 3 diversified across Trees and Bonds but likewise did not hedge because Tree holdings decreased in odd periods.

The subject-level differences reported in Table X are not exceptional. The contrast between individual choices and choices aggregated over individuals of the same type is stark.

In terms of choice predictions, the theory “works” at the type level but not at the individual level, which suggests that one has to be careful when extrapolating observations at the individual level to phenomena at the market level. If we had considered Subjects 3 or 5 as “representative,” and had predicted cross-sectional and temporal behavior of prices on the basis of their choices, the fit would have been poor. The situation is reminiscent of the cross-sectional variation in choices in static asset pricing experiments. There too, prices at the market level can be “right” (satisfy, for example, CAPM) even if individual choices are at odds with the theory. See Bossaerts, Plott, and Zame (2007).

Table X shows that subjects in certain periods held zero units of a security. If subjects followed equilibrium strategies and their risk aversion was not too high, this should not have happened; see Table BII in Appendix B. The following provides some idea of how often subjects traded to the boundary. We exclude period 1, because as mentioned before, it often took more than one period for subjects to trade away from initial holdings (all subjects started with zero holdings of one of the two securities). Of the 1,272 period-subjects (one subject per period), 19% held zero of the Tree and 24% held zero of the Bond. Are these numbers high? This is difficult to say. Again, holdings of individual subjects are enormously variable and this variability reveals nothing about the success of the Lucas model.

Table X HERE

Summary

Altogether, many facets of the Lucas model are confirmed in the experimental data. Relative price levels were correct, prices over time moved with fundamentals, and this movement was more pronounced when the Tree was cheaper. The Lucas model of asset pricing relies on Pareto optimality of consumption flows, and we find that average end-of-period cash holdings across subject types were indeed positively correlated, effectively undoing, through trading, the negative correlation of autarky consumptions.

Still, prices were excessively volatile, which made returns extremely noisy. While the equity premium was countercyclical, as predicted, its sign was actually incorrect in the high state. Evidently, the noise in prices translated into incorrect average returns.

Finally, while average consumptions across types revealed substantial Pareto improvement, individual consumption flows (end-of-period cash holdings) were highly variable, exhibiting no particular pattern. One would be hard-pressed to find a single individual who acted as in the theory, smoothing consumption across states and across odd and even periods, and hedging fundamental risk.

VI. Using Laboratory Data to Assess The Statistical Properties of GMM Tests of The Lucas Model

As noted before, we (control and) observe much more in our laboratory environment than is possible in the field. However, an interesting and potentially revealing exercise is to

ignore this additional richness, treat our laboratory data as if they were historical field data, and carry out the same kind of econometric tests that have been used for historical field data. Hence, let us consider only the time series of asset returns and aggregate consumption in our laboratory data. We ignore all additional information (true dividend process, true realized state, individual choices, etc). We then use GMM to test whether the stochastic Euler equations are satisfied. We assume, as has been done in analysis of historical field data, the presence of a representative agent with CRRA (power utility). We estimate the coefficient of risk aversion γ and the discount factor (β) and test whether the Euler equations hold, while imposing the constraint that the representative agent “consumes” the aggregate cash each period.

Note that the assumption of power utility is without loss of generality because we have only two states. Indeed, in equilibrium only the marginal rate of substitution μ between consumption in a low state and consumption in a high state is a free parameter (besides the time discount rate β). See Prediction 3. Power utility is one type of preference that could work; any other preference profile with decreasing marginal utility would fit too.

It is our opinion that the theory explains many facets of the experimental results. However, as mentioned before, the behavior of asset returns is anomalous along one important dimension. Specifically, Table VII shows that the average return on the Tree in the high state is significantly lower (at 2.0%) than that of the Bond (at 15.3%). Thus, the equity premium in the high state is (significantly) negative. The corresponding Euler equation can only be fit with a negative risk aversion coefficient (assuming that β is set at its correct

value of $5/6$). In contrast, in the low state, the ranking of returns is consistent with risk aversion, at 17.8% for the Tree and 16.1% for the Bond (though the difference is insignificant at $p = 0.05$). GMM will record the tension between the equity premium in the low state and in the high state, as long as one uses the (true) state as an instrument so that the conditional average returns across the two states are contrasted. The GMM statistic should then reject the model.

In the field, the true state of the economy (high, low) is rarely available. One immediately wonders whether traditional instruments would be good proxies. Our experimental data provide a unique opportunity to investigate this question. Since the GMM test should reject, what we are after is an analysis of the power of the GMM test when one uses the instruments that we know should lead to rejections, and when one uses proxies as instruments instead.

The Euler equations are

$$E \left[\beta \left(\frac{c_{t+1}^*}{c_t^*} \right)^{-\gamma} \frac{d_A + p_{A,t+1}}{p_{A,t}} - 1 \middle| I_t \right] = 0,$$

where c_t^* and c_{t+1}^* denote aggregate (per capita) consumption in periods t and $t + 1$, respectively, $A \in \{B, T\}$, and I_t is the information that agents in the economy (participants in our experiments) had at the end of period t . As is standard in GMM tests of these Euler equations, we choose variables z_t (“instruments”) in the agents’ information set. Each in-

strument generates a set of two unconditional moment conditions (one for each of the assets, B and T), by applying the law of iterated expectations:

$$\begin{aligned} & E \left[E \left[\left(\beta \left(\frac{c_{t+1}^*}{c_t^*} \right)^{-\gamma} \frac{d_A + p_{A,t+1}}{p_{A,t}} - 1 \right) z_t \middle| I_t \right] \right] \\ &= E \left[\left(\beta \left(\frac{c_{t+1}^*}{c_t^*} \right)^{-\gamma} \frac{d_A + p_{A,t+1}}{p_{A,t}} - 1 \right) z_t \right] \\ &= 0. \end{aligned}$$

Our first test is based on a traditional instrument choice, going back to Hansen and Singleton (1983). We choose as instruments (i) the constant one, (ii) lagged consumption growth, (iii) lagged returns on the Tree T , and (iv) lagged returns on the Bond B . Thus, we have four instruments and in turn, eight moment conditions. Only two parameters, β and γ , need to be estimated, so we have six overidentifying restrictions. The idea behind GMM is to find values of the parameters that minimize a quadratic form in the moment conditions. With a suitable weighting matrix, the resulting minimum is χ^2 distributed, with degrees of freedom equal to the number of overidentifying restrictions.³³ The necessary time series of consumption growth and asset returns are constructed by concatenating periods across all replications and all sessions, leaving out observations that would straddle two different replications, as we did for Table VI.

The top panel of Table XI displays the results of the first test. Three observations are worth noting.

1. The model is not rejected: $p = 0.310$.
2. The estimated discount factor β is significantly different from the theoretical one.
3. The coefficient of risk aversion γ is not significantly different from zero.

However, the results of this test may be misleading. In particular, GMM produces an estimate of the risk aversion coefficient that implies that the representative agent is risk-neutral or even slightly risk-loving, yet the price level data clearly show that subjects were risk averse, because the Tree is cheaper than the Bond. Moreover, participants smoothed consumption both across states and over time, again suggesting risk aversion. Further inspection of the data (Table VII) suggests why GMM produces this peculiar estimate of risk aversion. Average returns on the two securities, while positive, are little different from each other. The average equity premium is a tiny 0.1%, while there is a large difference in the covariation of dividends with the aggregate dividend. As a result, GMM fits parameters that make one believe that the representative agent is risk neutral.

Table XI HERE

In looking more closely, we might view the χ^2 GMM test of overidentifying restrictions as suspect. Two instruments – the lagged returns on the Tree and the Bond – are “weak,” in the sense that they are uncorrelated (even independent) over time, both with themselves and with consumption growth. Hence, these moment conditions do not provide additional restrictions beyond those imposed by the moment conditions constructed with the constant

as instrument. Effectively, the number of degrees of freedom in the χ^2 test is not six, but only two.

To determine the impact of these weak instruments, we run a second test, reestimating the model with only the constant and lagged consumption as instruments. The second row of Table XI displays the results. The model still fails to be rejected (with an even higher p -level). And the estimation of risk attitude continues to be anomalous, because the point estimate remains in the risk-loving range, though insignificantly different from risk neutrality.

Poor instrument choice may be the reason behind the low power of GMM and its tendency to generate the wrong estimates of risk aversion. As argued before, anomalies emerge specifically when contrasting average returns across the high and low states. Lagged aggregate consumption growth is a noisy proxy for the state. What if we used the true state as an instrument? Would we obtain the rejection that we expect? And would the estimate of the coefficient of risk aversion become significantly positive, in line with the choices of our subjects and with the patterns we observe in price levels?

Consequently, in our third test, we replace the constant instrument with two dummy variables, one that tracks the high state and another one that tracks the low state. We keep the remaining instrument, consumption growth. This gives three instruments and as such generates six moment conditions. With two parameters to estimate, we are left with four degrees of freedom. The results are presented in the bottom panel of Table XI. We observe the following.

1. The model is now rejected ($p < 0.01$).
2. The discount factor, β , remains a bit too high.
3. Risk aversion is now highly significant ($p = 0.001$).

In sum, when the true state is used as an instrument, we recover power (to reject) and obtain parameter estimates that are consistent with subjects' choices as well as with patterns in price levels.

The impact of instrument choice on the power of the GMM test should caution anyone when conclusions need to be drawn from field data. Perhaps most annoying is the misestimation of the coefficient of risk aversion when using traditional instruments. Because the choice of utility function in the estimation is without loss of generality, it would be wrong to search for alternative preferences just because we estimated risk aversion to be too low. The misestimation is solely due to lack of power when implementing the GMM test with traditional instruments.

VII. The Expected and The Anomalous

Viewed against the predictions of the Lucas model, our experiments generate findings that are expected – prices and individual consumption are correlated with fundamentals (aggregate consumption) and prices are ordered across both securities and states as predicted by the theory – and findings that seem anomalous – prices are excessively volatile and

subjects do not hedge. Because volatile prices would seem to signal clearly the need to hedge, the co-existence of excess volatility and lack of hedging seems surprising. However, the particular kind of excess volatility that we see in the experimental data might well lead subjects to conclude that there is no opportunity to hedge.

To see why this might be so, recall that the predictions of the Lucas model, and indeed the very definition of Radner equilibrium, depend on the assumption that agents have perfect foresight. This does not mean that agents can predict the random future dividends, but rather that they know the random process of dividends and what prices will obtain for each possible dividend realization. In the experiment, the subjects are told the dividend process, so in principle any lack of perfect foresight cannot be there. But the price process must be learned. It seems to be too much to expect that price beliefs be exactly correct. Perhaps it is more realistic to expect that beliefs be *approximately* correct?

Optimization against exactly correct beliefs leads exactly to the Radner equilibrium predicted by the Lucas model. It would seem that optimization against approximately correct beliefs should lead to something that approximates the Radner equilibrium predicted by the Lucas model. However, this is not necessarily so: because the price process is *endogenous*, beliefs about the price process can be approximately correct (i.e., close to the realized price process) and still very far from the price process predicted by the Lucas model. The same point is made by Adam, Marcet, and Nicolini (2015), who use it to explain excess volatility in historical data.

On the basis of our experimental data, it seems quite plausible that agents expected prices to follow a martingale, as would be predicted by traditional accounts of EMH, and *not* to co-move with economic fundamentals, as would be predicted by the Lucas model. In Section V, we document how subjects did not hedge price risk, in contrast with optimal trading strategies in the Lucas equilibrium. Hedging is accomplished by buying Trees when subjects experience an income shortfall. In equilibrium, Trees are a perfect hedge against future price changes, when income flows allow subjects to repurchase securities: Trees pay dividends when prices are high but not when prices are low. If subjects decided not to hedge price risk by buying Trees, and if their decisions truly reflect their beliefs, then subjects appeared to believe that prices move in ways that are unrelated to dividends. That is, choices revealed expectations according to which prices are a martingale (possibly with drift, but definitely unrelated to fundamentals).

The belief that prices follow a martingale is wrong, but it is not readily falsifiable on the basis of the limited number of observations available to subjects. Indeed, the belief that, for example, Bond prices follow a martingale would be hard to falsify even after 80 observations (see Table VI), which is an order of magnitude more observations than were available to subjects. The martingale belief is therefore a credible working hypothesis.

A thought experiment may help one understand the consequences of incorrect beliefs. Imagine that in every period agents always believe that past prices are the best predictions of future prices, independent of economic fundamentals; that, given these beliefs, agents correctly solve their current optimal investment-consumption problem as a function of prices;

that agents then send demand schedules to the market; and that markets generate prices so that demand and supply are equal in that period. Of course, beliefs will be revealed to be wrong next period, so in this thought experiment, we are considering a kind of temporary equilibrium (Grandmont (1977)) in which beliefs, although incorrect, are disciplined by observation. How would prices in this temporary equilibrium evolve over time? Simulations suggest that prices would evolve very much as in the experiment: they co-move with dividends, but very noisily – hence, they would be excessively volatile.

To illustrate this thought experiment, we carry out simulations. There are two types of agents, endowed as in the experiment, and each type is represented by an agent with logarithmic utility. Agent beliefs (that prices revert to the levels of the previous period) are affected every period by an additive Gaussian disturbance with mean zero and standard deviation \$0.40. Agents start out believing that the Tree will be priced at \$2.5 and the Bond at \$3. Figure 3 displays the evolution of prices and states in a typical simulation. Price dynamics are in line with those in the experiment – prices relate to states only in a noisy way. At the same time, agents turn out not to hedge residual price forecasting risk; they accommodate income shortfalls solely by selling Bonds and Trees. Still, their choices move substantially towards Pareto optimality: the consumption share of the Type I agent fluctuates between 39% and 44%, little affected by state and period (odd/even).

Figure 3 HERE

This thought experiment demonstrates starkly that the price predictions of the Lucas model are sensitive to small mistakes in beliefs about the price processes. This comes as a surprise because price predictions in the Lucas model are robust to small mistakes in beliefs about the dividend processes (Hassan and Mertens (2010)). Consequently, it seems that mistakes in beliefs can manifest themselves quite differently: when they concern the dividend processes, mistakes are dampened out; when they concern price processes, mistakes can create positive feedback and hence are amplified.

VIII. Conclusion

Over the last 30 years, the Lucas model of dynamic asset pricing has become the lens through which scholars of finance and macroeconomics interpret historical data coming from the field. Yet little is known about the true relevance of the Lucas model, and confidence in the model has been shaken by recent events. This paper was prompted in part by the belief that proper understanding of the Lucas model – and of the thinking underlying it and the applications that are made of it – could be greatly advanced if we could examine the workings of a Lucas economy in the laboratory. Of course, it is a long way from the laboratory to the real world. Many features of the real world are absent in the laboratory and these features may well have an enormous bearing on the applicability of the Lucas model (or any other model). But this seems to us to argue even more forcefully for laboratory experimentation.

Models are idealizations and the laboratory is an idealized environment; if the models do not work in the laboratory, why should we expect them to work in the real world?

In our view, our experiments provide substantial support for the Lucas asset pricing model. Our experimental results display a number of features that are consistent with the most important predictions of the Lucas model: prices move with fundamentals, agents trade assets to smooth consumption and insure against risk, more risky assets yield a substantial premium over less risky assets, and the equity premium is countercyclical. However, our experimental results do display a feature that is at odds with the predictions of the Lucas model: prices display excess volatility, so returns are noisy,³⁴ to the point that the equity premium is actually negative in “booms.” This causes rejections in standard GMM tests of the stochastic Euler equations (provided the right instruments are employed). Interestingly, excess volatility is precisely the feature of field data that much of the literature has attempted to explain in terms of “frictions” or deviations from the basic model. These frictions and deviations are entirely absent in our laboratory environment, suggesting that they may not be needed to explain the empirical failure of the model.

Instead, we conjecture that excessive volatility was caused by subjects’ mistakes in forecasting future prices. The Lucas model presumes that the agents perfectly forecast future asset prices, but the subjects in our experiments could not and did not do so. Indeed, the absence of a hedging component to choices reveals that subjects did not expect prices to correlate with fundamentals. When we simulated the temporary equilibrium (Grandmont (1977)) of our laboratory economy, we found that beliefs that prices would not change with

fundamentals were difficult to falsify – such beliefs generated “almost” perfect foresight. At the same time, prices were far more volatile than in the perfect-foresight equilibrium of the Lucas model. Interestingly, little Pareto efficiency was lost because, at the end of the day, agents in the simulations made only small forecast mistakes. In that dimension too, our experiments are in line with our conjecture: we recorded substantial Pareto improvements.

It might be argued that, because our experiments are short (the longest replication is only 10 periods) we have not provided a fair opportunity for subjects to learn perfect forecasts. However, in many ways we gave our subjects a much simpler problem than they face in the real world. In our experiment, subjects are told the true dividend process, whereas in the field they would have to learn it. Further, in our experiment the true dividend process *is stationary*, while in the field it is not. (Indeed, it might be argued that the world is not stationary at all, and that stationarity is just an assumption imposed on a model that would otherwise be completely intractable.) Thus, in many dimensions, we gave the Lucas model the best possible chance to succeed.³⁵

It would be of interest to know whether our subjects would eventually learn to make correct forecasts if they had many more observations. Unfortunately, a design that provides enough observations seems quite impractical. We have noted that it would require 80 periods for subjects to learn about the true relation between prices and fundamentals; to obtain replications of 80 periods would require that we choose the continuation probability β uncomfortably close to one. A high continuation probability may induce subjects not to take seriously the trade-off between cash on hand and securities holdings. Perhaps more

importantly, a replication of 80 periods would last at least four to five hours (following the very substantial initial time for training); carrying out an experimental session of such a duration would be difficult.

If agents are aware that, besides fundamental risk, they face risk because their foresight is imperfect, then the equilibrium (Euler) equations will be misleading. This is because the equations assume that there are no residual price forecasting mistakes. One would expect the agents themselves to be at least as averse to fundamental risk (about which they know a lot) as to residual price forecasting mistakes (about which they do not know much). Their investments, and hence consumption choices, are likely to be affected. But the equilibrium equations explain consumption choices only in terms of fundamental risk, and consequently there is a real danger that agents will appear far more risk averse through those equations than they actually are. Future experiments should clarify the extent to which awareness of residual price forecasting mistakes introduce risk premia that are missing in the perfect foresight equilibrium on which the Lucas model is built.

Our experimental findings also illustrate that one should be cautious when extrapolating from the individual to the market. As in our static experiments (Bossaerts, Plott, and Zame (2007)), we find substantial heterogeneity in choices across subjects; most individual choices have little or no explanatory power for market prices, or even for choices averaged across subjects of the same type (same endowments). Overall, the system (market) behaves as predicted by the theory (at least qualitatively), but individual choices do not. Hence, we caution strongly against giving too much credence to asset pricing theories in which the

system is simply a mirror image of one of its parts. The “laws” of the (financial) system may be different from those of its parts.

The idea of looking at experimental findings using the methodology typically used to study historical field data is borrowed from Asparouhova (2006). We think that the exercise yields interesting insights. We find that statistical inference is biased unless one has access to the “correct” instruments. Outside the laboratory, it would be hard to determine what the correct instruments are, and even if one knew, they may not be observable.

References

- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini, 2015, Stock market volatility and learning, *Journal of Finance*, forthcoming.
- Arrow, Kenneth J., Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz, 2008, The promise of prediction markets, *Science* 320, 877-878.
- Asparouhova, Elena, 2006, Competition in lending: Theory and experiments, *Review of Finance* 10, 189-219.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

- Bernard, Victor L., and Jacob K. Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research* 27, pp. 1-36.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, pp. 637-654.
- Bossaerts, Peter, Debrah Meloso, and William Zame, 2013, Dynamically-complete experimental asset markets, Working paper.
- Bossaerts, Peter, Charles Plott, and William Zame, 2007, Prices and portfolio choices in financial markets: Theory, econometrics, experiment, *Econometrica* 75, 993-1038.
- Bossaerts, Peter, and William R. Zame, 2006, Asset trading volume in infinite-horizon economies with dynamically complete markets and heterogeneous agents: Comment, *Finance Research Letters* 3, 96-101.
- Bossaerts, Peter, and William Zame, 2008, Risk aversion in laboratory asset markets, in *Research in Experimental Economics* 12, JAI Press, Greenwich, CT.
- Bruguier, Antoine J., Steven R. Quartz, and Peter Bossaerts, 2010, Exploring the nature of “trader intuition,” *Journal of Finance* 65, 1703-1723.
- Camerer, Colin F., and Keith Weigelt, 1996, An asset market test of a mechanism for inducing stochastic horizons in experiments, in *Research in Experimental Economics* 6, JAI Press, Greenwich, CT.

- Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205-251.
- Cochrane, John H., 2001, *Asset Pricing*, Princeton University Press, Princeton, NJ.
- Crockett, Sean, and John Duffy, 2013, An experimental test of the Lucas asset pricing model, Working paper, IDEAS.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839-1885.
- De Bondt, Werner F. M., and Richard Thaler, 1985, Does the stock market overreact? *Journal of Finance* 40, 793-805.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379-395.
- Duffie, Darrell, and Chi-Fu Huang, 1985, Implementing Arrow-Debreu equilibria by continuous trading of few long-lived securities, *Econometrica* 53, 1337-1356.
- Dunn, Kenneth B., and Kenneth J. Singleton, 1986, Modeling the term structure of interest rates under non-separable utility and durability of goods, *Journal of Financial Economics* 17, 27-55.
- Epstein, Larry G., and Tan Wang, 1994, Intertemporal asset pricing under Knightian uncertainty, *Econometrica* 62, pp. 283-322.

- Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, 263-286.
- Geanakoplos, John, and William R. Zame, 2014, Collateral equilibrium, I: A basic framework, *Economic Theory* 56, 443-492.
- Goyal, Amit, and Ivo Welch, 2003, Predicting the equity premium with dividend ratios, *Management Science* 49, 639-654.
- Grandmont, Jean Michel, 1977, Temporary general equilibrium theory, *Econometrica* 45, 535-572.
- Hansen, Lars, and Kenneth J. Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy* 91, 249-265.
- Hassan, Tarek A., and Thomas M. Mertens, 2011, The social cost of near-rational investment: Why we should worry about volatile stock markets, Working paper 17027, National Bureau of Economic Research.
- Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163-1198.
- Hernandez, Alejandro, and Manuel S. Santos, 1996, Competitive equilibria for infinite-horizon economies with incomplete markets, *Journal of Economic Theory* 71, 102-130.
- Holt, Charles, and Susan Laury, 2002, Risk aversion and incentive effects, *American Economic Review* 92, 1644-1655.

- Judd, Kenneth L., Felix Kubler, and Karl Schmedders, 2003, Asset trading volume with dynamically complete markets and heterogeneous agents, *Journal of Finance* 58, 2203-2218.
- Kehoe, Timothy J., and David K. Levine, 1993, Debt-constrained asset markets, *Review of Economic Studies* 60, 865-888.
- LeRoy, Stephen F., and Richard D. Porter, 1981, The present-value relation: Tests based on implied variance bounds, *Econometrica* 49, 555-574.
- Levine, David K., and William R. Zame, 1996, Debt constraints and equilibrium in infinite horizon economies with incomplete markets, *Journal of Mathematical Economics* 26, 103-131.
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41-66.
- Lucas, Robert E. Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429-1445.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *Journal of Finance* 60, 1167-1219.
- Maenhout, Pascal J., 2004, Robust portfolio rules and asset pricing, *Review of Financial Studies* 17, 951-983.

- Magill, Michael, and Martine Quinzii, 1994, Infinite horizon incomplete markets, *Econometrica* 62, 853-880.
- Malkiel, Burton G., 1999, *A Random Walk Down Wall Street*, W.W. Norton, New York, NY.
- Martin, Ian, 2011, The lucas orchard, *Econometrica* 81, 55-111.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145-161.
- Merton, Robert C., 1973a, An intertemporal capital asset pricing model, *Econometrica* 41, 867-887.
- Merton, Robert C., 1973b, Theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141-183.
- Nagel, Stefan, and Kenneth J. Singleton, 2011, Estimation and evaluation of conditional asset pricing models, *Journal of Finance* 66, 873-909.
- Plott, Charles, and Shyam Sunder, 1988, Rational expectations and the aggregation of diverse information in laboratory security markets, *Econometrica* 56, 1085-1118.
- Rabin, Matthew, 2000, Risk aversion and expected-utility theory: A calibration theorem, *Econometrica* 68, 1281-1292.
- Radner, Roy, 1972, Existence of equilibrium of plans, prices, and price expectations in a sequence of markets, *Econometrica* 40, 289-303.

Routledge, Bryan R., and Stanley E. Zin, 2010, Generalized disappointment aversion and asset prices, *Journal of Finance* 65, 1303-1332.

Samuelson, Paul A., 1973, Proof that properly discounted present values of assets vibrate randomly, *Bell Journal of Economics and Management Science* 4, 369-374.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421-436.

Smith, Vernon L., 1965, Experimental auction markets and the Walrasian hypothesis, *Journal of Political Economy* 73, 387-393.

Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams, 1988, Bubbles, crashes, and endogenous expectations in experimental spot asset markets, *Econometrica* 56, 1119-1151.

Appendix A: Theory

The main text only discusses theoretical predictions that have empirical relevance, leaving out details of the arguments that support those predictions. Here, we list the main theoretical results in a logically coherent way, which necessarily implies some repetition of points already made in the text.

1. *Individual consumption is stationary and perfectly correlated with aggregate consumption.*

To see this, fix a period t and a state $\sigma = H, L$. The boundary condition guarantees that equilibrium allocations are interior, so Pareto optimality guarantees that all agents have the same marginal rate of substitution for consumption in state σ at periods t and $t + 1$: $u'_i(c_{i,t+1}^\sigma)/u'_i(c_{i,t}^\sigma) = u'_j(c_{j,t+1}^\sigma)/u'_j(c_{j,t}^\sigma)$ for each i, j . In particular, the ranking of marginal utility for consumption in state σ at dates t and $t + 1$ must be the same for all agents. Because utility functions are strictly concave, the rankings of consumption in state σ at dates t and $t + 1$ must be the same for all agents (and opposite to rankings of marginal utilities). But the sum of individual consumptions is aggregate consumption, which is stationary, and hence equal in state σ at periods t and $t + 1$. Hence the consumption of each individual agent must also be equal in state σ at periods t and $t + 1$. Since t is arbitrary, this means that individual consumption must be constant in state σ , that is, stationary. Because the rankings of consumption across states are the same

for all agents, the ranking must agree with the ranking of aggregate consumption, so individual consumption is perfectly correlated with aggregate consumption.

2. *The stochastic Euler equations obtain both at the individual level and for the representative consumer.*

To see this, fix an agent i , and write $\{c_i\}$ for i 's stochastic equilibrium consumption stream (which we have just shown to be stationary). Because i optimizes given current and future asset prices, asset prices in period t must equalize marginal utility of consumption at each state in period t with expected marginal utility of consumption at period $t + 1$. If i buys (sells) an additional infinitesimal amount ε of asset $A = B, T$ at period t , consumption in period t is reduced (increased) by ε times the price of the asset but consumption in period $t + 1$ is increased (reduced) by ε times the delivery of the asset, which is the sum of its dividend and its price in period $t + 1$. Hence, the first-order condition is

$$p_{A,t}^\sigma = \beta \left\{ \pi \left[\frac{u'_i(c_i^H)}{u'_i(c_i^\sigma)} \right] (d_A^H + p_{A,t+1}^H) + (1 - \pi) \left[\frac{u'_i(c_i^L)}{u'_i(c_i^\sigma)} \right] (d_A^L + p_{A,t+1}^L) \right\},$$

where superscripts index states and subscripts index assets, time, and agents in the obvious way. We can write this in more compact form as

$$p_{A,t}^\sigma = \beta E \left\{ \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} \right] (d_A + p_{A,t+1}) \right\} \quad (A1)$$

for $\sigma = H, L$ and $A = B, T$. Equation (A1) is the familiar stochastic Euler equation except that the marginal utilities are those of an arbitrary agent i and not of the representative agent. (Equality of the ratios of marginal utilities across agents, which is a consequence of Pareto optimality, implies that (A1) is independent of the choice of agent i and that we could write (A1) in terms of the utility function of a representative agent – but the utility function of the representative agent would be determined *in equilibrium*.)

3. *Asset prices are stationary.*

Fix an asset $A = B, T$ and a period t . The stochastic Euler equation (A1) expresses prices $p_{A,t}$ at time t in terms of marginal rates of substitution, dividends, and prices at times $t + 1$. Substituting $t + 1$ for t expresses prices $p_{A,t+1}$ at time $t + 1$ in terms of marginal rates of substitution, dividends, and prices at times $t + 2$, and so forth. Combining all these substitutions and keeping in mind that consumptions, marginal rates of substitution, and dividends are stationary yields an infinite series for prices

$$\begin{aligned}
 p_{A,t}^\sigma &= \sum_{\tau=0}^{\infty} \beta^{\tau+1} E \left[\frac{u'_i(c_{i,t+\tau+1})}{u'_i(c_{i,t}^\sigma)} d_{A,t+\tau+1} \right] \\
 &= \beta E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right] \sum_{\tau=0}^{\infty} \beta^\tau \\
 &= \left(\frac{\beta}{1-\beta} \right) E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right]. \tag{A2}
 \end{aligned}$$

The terms in the infinite series are stationary so prices are stationary as well.

4. *Asset prices are determined by one unknown parameter.*

Let $\mu = u'_i(c_i^L)/u'_i(c_i^H)$ be the marginal rate of substitution in the low state for consumption in the high state (which Pareto optimality guarantees is independent of which agent i we use); note that risk aversion implies $\mu > 1$. The assertion then follows immediately from (A2) but a slightly different argument is perhaps more revealing. For each asset $A = B, T$, we can write the stochastic Euler equations as

$$\begin{aligned} p_A^H &= \beta \left[\pi(d_A^H + p_A^H) + (1 - \pi)(d_A^L + p_A^L)\mu \right] \\ p_A^L &= \beta \left[\pi(d_A^H + p_A^H)(1/\mu) + (1 - \pi)(d_A^L + p_A^L) \right]. \end{aligned}$$

It follows immediately that

$$p_A^H/p_A^L = \mu. \tag{A3}$$

Substituting and solving yields

$$\begin{aligned} p_A^H &= \left(\frac{\beta}{1 - \beta} \right) \left[\pi d_A^H + (1 - \pi)d_A^L \mu \right] \\ p_A^L &= \left(\frac{\beta}{1 - \beta} \right) \left[\pi d_A^H (1/\mu) + (1 - \pi)d_A^L \right]. \end{aligned} \tag{A4}$$

Specializing to the parameters of the experiment $d_T^H = 1, d_T^L = 0, d_B^H = d_B^L = 0.5$, and

$\beta = 5/6$ yields

$$p_B^H = (2.5)(1 + \mu)/2$$

$$p_B^L = (2.5)(1 + \mu)/2\mu$$

$$p_T^H = 2.5$$

$$p_T^L = 2.5/\mu.$$

Only one parameter, μ , determines prices of both securities in both states. Note that

$p_T^H = 2.5$ (the price of the tree in the high state is independent of risk attitudes) and

$p_B^H/p_B^L = p_T^H/p_T^L$ (the ratios of asset prices in the two states are the same).

5. *Asset prices are perfectly correlated with fundamentals.*

This is also an immediate consequence of equations (A4), because $\mu > 1$ asset prices are higher in the high state than in the low state. Informally, this is understood most clearly by thinking about the representative agent. In state H , aggregate consumption supply is high, so high prices (low returns) must be in place to temper the representative agent's desire to save (buy). The opposite is true for state L : aggregate consumption is low, so low prices (high returns) temper the representative agent's desire to borrow (sell).

6. *The Tree is cheaper than the Bond.*

This too is a consequence of equations (A4). In the context of static asset pricing

theory this pricing relation is a simple consequence of the fact that the dividends on the Tree have higher covariance with aggregate consumption than does the Bond; the Tree has higher “beta” than the Bond. However, in the dynamic context the result is more subtle because asset prices in period t depend on dividends in period $t + 1$ *and* on asset prices in period $t + 1$; since prices are determined in equilibrium, it is not automatic a priori that *prices* of the Tree have higher covariance with aggregate consumption than prices of the Bond.

7. *Expected returns on the Tree and Bond change with the state and are countercyclical.*

From prices and dividends of the Tree and Bond as well as the state transition probabilities (from a low state one moves to a high state or remains in the low state with equal probability), one can readily compute the expected returns on the Tree and the Bond. Simple algebraic manipulation then allows one to express the difference across high and low states of the expected return on the Tree ($E[R_T|H] - E[R_T|L]$) and Bond ($E[R_B|H] - E[R_B|L]$) as follows:

$$\begin{aligned} E[R_T|H] - E[R_T|L] &= \pi(1 - \mu) + (1 - \pi)\left(\frac{1}{\mu} - 1\right) + d_B(1 - \mu)\frac{1}{p_T^H}, \\ E[R_B|H] - E[R_B|L] &= \pi(1 - \mu) + (1 - \pi)\left(\frac{1}{\mu} - 1\right) + d_B(1 - \mu)\frac{1}{p_B^H}. \end{aligned}$$

Because $\mu > 1$, all terms in both expressions are negative, and hence the expected return on both assets is higher in “busts” (when the state is low) than in “booms” (when the state is high).

8. *The equity premium is positive and countercyclical.*

The difference between the expected return on the risky security (the Tree) and the expected return on the (relatively) risk-free security (the Bond) is the *equity premium* (Mehra and Prescott (1985)).³⁶ The conclusion that the Tree is cheaper than the Bond implies that the equity premium is positive. Because asset prices are stationary, equity premia are stationary as well; simple computations show that the equity premia in the high and low states are (remember that the expected dividends are the same for both assets and equal to d_B , and that for each asset A , $p_A^H = \mu p_A^L$):

$$\begin{aligned} E^H &= \frac{\pi p_T^H + (1 - \pi)p_T^L + d_B - p_T^H}{p_T^H} - \frac{\pi p_B^H + (1 - \pi)p_B^L + d_B - p_B^H}{p_B^H} \\ &= d_B \left(\frac{1}{p_T^H} - \frac{1}{p_B^H} \right) \\ E^L &= \frac{\pi p_T^H + (1 - \pi)p_T^L + d_B - p_T^L}{p_T^L} - \frac{\pi p_B^H + (1 - \pi)p_B^L + d_B - p_B^L}{p_B^L} \\ &= d_B \left(\frac{1}{p_T^L} - \frac{1}{p_B^L} \right) = d_B \left(\frac{\mu}{p_T^H} - \frac{\mu}{p_B^H} \right) = \mu E^H. \end{aligned}$$

Note that both equity premia are positive. The difference across states is

$$E^H - E^L = (1 - \mu)E^H.$$

This difference is strictly negative (because $\mu > 1$) so the equity premium is countercyclical (lower in the high state than in the low state). Note that countercyclicity

provides the correct incentives: when dividends are low, the equity premium is high, so investors buy risky Trees rather than consuming scarce dividends; when dividends are high, the equity premium is low, so investors prefer to consume rather than engage in risky investment.

9. *Cross-sectional and time-series properties of asset prices reinforce each other.*

To be more precise, as the discount of the Tree price relative to the Bond price increases because risk aversion rises, the difference in Tree prices or in Bond prices across states increases. That is,

$$\text{cov}(p_B^\sigma - p_T^\sigma, p_A^H - p_A^L) > 0,$$

for $\sigma = H, L$ and $A = B, T$, with covariance computed based on sampling across cohorts of agents (economies), keeping everything else constant. “Everything else” means: initial endowments, private income flows, asset structure, outcome probabilities, as well as impatience β . Economies are therefore distinguishable at the price level only in terms of the risk aversion (embedded in x) of the representative agent.³⁷

10. *Agents smooth consumption over time.*

Individual equilibrium consumptions are stationary but individual endowments are not, so agents smooth over time.

11. *Agents trade to hedge fundamental risk.*

If prices were constant, agents could smooth consumption simply by buying or selling one asset. However, there *is* price risk, because prices move with fundamentals and

fundamentals are uncertain. Hence, when agents sell assets because private income is low (relative to average private income), they also need to insure against the risk that prices might change by the time they are ready to buy back the assets. In equilibrium, prices increase with the dividend on the Tree, and agents correctly anticipate this. Since the Tree pays a dividend when prices are high, it is the perfect asset to hedge price risk. Consequently, but perhaps counterintuitively, agents *buy* Trees in periods when private income is low and *sell* when private income is high.

Appendix B: Numerical Example

Here, we provide explicit solutions for equilibrium prices, holdings, and consumptions when taking the parameters as in the experiment and assuming that all agents display identical CRRA, which we vary from 0.2 to 1.0.

- There are an even number $n = 2m$ of agents; agents $i = 1, \dots, m$ are of Type I, while agents $i = m + 1, \dots, 2m$ are of Type II.
- Type I agents are endowed with asset holdings $b_I = 0, \tau_I = 10$ and have income $e_{I,t} = 15$ when t is even and $e_{I,t} = 0$ when t is odd.
- Type II agents are endowed with asset holdings $b_{II} = 10, \tau_{II} = 0$ and have income $e_{II,t} = 15$ when t is odd and $e_{II,t} = 0$ when t is even.
- All agents have CRRA $\gamma = 0.2, 0.5, 1$. (There is nothing special about these particular choices of risk aversion; we offer them solely for comparison purposes. We note that

risk aversion in the range $0.2 - 0.5$ is consistent with the experimental findings of Holt and Laury (2002) and Bossaerts, Plott, and Zame (2008).)

- The initial state is high.

Table BI HERE

Table BI provides equilibrium asset prices, the discounts in the price of the Tree relative to the Bond, and equity premia as functions of the state and of risk aversion. As expected, Trees are always cheaper than Bonds. The discount on the Tree is higher in state H than in state L , while the equity premium is lower in state H than in state L , reflecting the procyclical behavior of the discount and the countercyclical behavior of the equity premium. The dependence of prices on the state and the predictability of returns are apparent from the table.³⁸

Table BII HERE

Table BII displays equilibrium holdings and trades for Type I agents, who receive income in even periods and face an income shortfall in odd periods. (Equilibrium holdings and trades of Type II agents are of course complements to those of Type I agents.) As expected, the absence of income in odd periods is resolved not through outright sales of assets, but through a *combination* of sales of Bonds and purchases of Trees. The Bond sales provide income; the Tree purchases hedge price risk over time.^{39,40}

Equilibrium holdings and trades ensure that Type I agents consume a *constant fraction* (48%) of total available consumption in the economy, independent of state or date; of course Type II agents consume the complementary fraction (52%). That consumption shares are constant is a consequence of the assumptions that allocations are Pareto optimal and that agents have identical homothetic utilities; as we note earlier, without the assumption of identical homothetic utilities all we can conclude is that individual consumptions are perfectly correlated with aggregate consumption.

Appendix C: Instruction Set

1. *Situation*

One session of the experiment consists of a number of replications of the same situation, referred to as periods. You will be allocated securities that you can carry through all periods.

You will also be given cash, but cash will not carry over from one period to another.

Every period, markets open and you will be free to trade your securities. You buy securities with cash and you get cash if you sell securities.

Cash is not carried over across periods, but there will be two sources of fresh cash in a new period. First, the securities you are holding at the end of the previous period may pay dividends. These dividends become cash for the subsequent period. Second, before the start of specific periods, you may be given income. This income becomes cash for the period. It will be known beforehand in which periods you receive income.

Each period lasts 5 minutes. The total number of periods is not known beforehand. Instead, at the end of a period, we determine whether the experiment continues, as follows. We throw a twelve-sided die. If the outcome is 7 or 8, we terminate the session. Otherwise we continue and advance to the next period. Notice: the termination chance is time-invariant; it does not depend on how long the experiment has been going.

Your experiment earnings are determined by the cash you are holding at the end of the period in which the session ends.

So, if you end a period without cash, and we terminate the session at that point, you will not earn any money for the session. This does not mean, however, that you should ensure that you always end with only cash and no securities. For in that case, if we continue the experiment, you will not receive dividends, and hence, you start the subsequent period without cash (and no securities) unless this is a period when you receive income.

We will run as many sessions as can be fit in the allotted time of two hours for the experiment. If the last session we run has not been terminated before the scheduled end of the experiment we will terminate the session and you will earn the cash you are holding at that point.

You will be paid the earnings of two randomly chosen sessions. If we manage to run only one session during the allotted time for the experiment, you will be paid double the earnings for that session. During the experiment, accounting is done in real dollars.

2. Data

There will be two types of securities, called tree and bond. One unit of the tree pays a random dividend of zero or one dollar, with equal chance; past dividends have no influence on this chance. (The actual draw is obtained using the standard pseudo-random number generator in the program matlab). One unit of the bond always pays fifty cents. You will receive the dividends on your holdings of trees and bonds in cash before a new period starts. As such, you will receive dividends on your initial allocation of trees and bonds before the first period starts.

You will start this session with 10 trees and 0 bonds. Others may start with different initial allocations. In addition, you will receive income every alternate period. In odd periods (1, 3,...) you will receive nothing, and in even periods (2, 4,...) you will receive 15 dollars. This income is added to your cash at the beginning of a new period. Others may have a different income flow.

Because cash is taken away at the end of a period when the session does not terminate, the dividend payments you receive, together with your income, are the sole sources of cash for a new period.

Figure C1 displays the timeline of events during one replication.

Figure C1 HERE

3. Examples (for illustration only)

Figure C2 displays two sample tables of outcomes in a session. It is assumed that the session ends after the 6th period. The top panel in Figure C2 shows a table with asset holdings, dividend and cash each period if the states are as per row 2 and the individual sticks to the initial allocation throughout. The final take-away cash/earning is 25 dollars as the session terminated after 6th period. It would have been \$0 if it had terminated in period 5.

The bottom panel of Figure C2 shows the case where the individual trades as follows:

- In period 1, to an allocation of 5 trees and 5 bonds.
- And subsequently, selling to acquire more cash if dividends and income are deemed too low.
- Or buying more assets when dividends and income are high.

Since there is a $1/6$ chance that the session ends in the period when a security is bought, its expected value equals $(5/6)/(1-5/6)$ times the expected dividend (which is equal for both the tree and the bond), or $5 \cdot 0.5 = 2.5$. Trade is assumed to take place at 2.5. Note, however, that the actual trading prices may be different, and that they may even change over time, depending on, e.g., the dividend on the tree. The final take-away cash in this case is \$15 dollars. It would have been \$13 if the session had terminated in period 5.

Figure C2 HERE

Tables

Table I
Summary Data

This table presents summary data for all experimental sessions. Stars indicate sessions when last replication ended with a switch to the two-period termination protocol.

Session	Place	Number of Replications	Number of Periods (Total within Session, Min. across Replications, Maximum)	Subject Count
1	Caltech*	4	(14, 1, 7)	16
2	Caltech	2	(13, 4, 9)	12
3	UCLA*	3	(12, 3, 6)	30
4	UCLA*	2	(14, 6, 8)	24
5	Caltech*	2	(12, 2, 10)	20
6	Utah*	2	(15, 6, 9)	24
Overall		15	(80, 1, 10)	126

Table II
Trading Volume

This table presents the trading volumes of the two traded securities for all experimental sessions. The first panel aggregates the trading volume data across all periods. The second panel presents the data for periods 1 and 2 only. The third panel aggregates the data from period 3 on.

Periods	Security	Mean	St. Dev.	Min	Max
All	Tree	23	12	3	59
	Bond	17	11	2	58
1 and 2	Tree	30	15	5	59
	Bond	21	14	4	58
≥ 3	Tree	19	8	3	36
	Bond	15	9	2	41

Table III
Summary Statistics of Transaction Prices

This table presents summary statistics on transaction prices for the two traded securities. The last row displays the corresponding discount of the price of the Tree relative to that of the Bond.

	Mean	St. Dev.	Min	Max
Tree	2.75	0.41	1.86	3.70
Bond	3.25	0.49	2.29	4.32
Discount	0.50	0.40	-0.20	1.79

Table IV
Period-Average Transaction Prices per State

This table presents period-average transaction prices for the two traded securities across the high and low states. The last row displays the corresponding discount of the Tree price relative to the Bond price.

	High State	Low State	Difference
Tree	2.91	2.66	0.24
Bond	3.34	3.20	0.14
Discount	0.43	0.54	-0.11

Table V
Discount and Cross-Sectional Price Difference Correlation Across Replications
This table presents the correlations between (i) the average discounts, per replication, on the Tree price relative to the Bond price, and (ii) the average differential between high and low states of the Tree or Bond prices.

	Tree	Bond
Correlation	0.80	0.52
(St. Err.)	(0.40)	(0.40)

Table VI
OLS Regression of Changes in Period-Average Transaction Prices

This table presents the results of an OLS regression analysis of changes in period-average transaction prices. * indicates that corresponding statistics is significant at $p = 0.05$.

Explanatory Variables	Tree Price Change		Bond Price Change	
	Estimate	95% Conf. Int.	Estimate	95% Conf. Int.
Change in State Dummy (None=0; High-to-Low=-1, Low-to-High=+1)	0.19*	(0.08, 0.29)	0.10	(-0.03, 0.23)
R^2		0.18		0.04
Autocor. (s.e.=0.13)		0.18		-0.19

Table VII**Average Returns Across Securities and States**

This table presents the average returns across securities and across states. The last row displays the equity premium. Standard errors in parentheses. ** indicates one-sided significance at the 1% level (Equity Premium: t test of difference in return across securities; Difference across states: two-sample unequal-variance t test).

	All Periods	High State Periods	Low State Periods	Difference High/Low
Tree	12.8 (%) (3.3)	2.0 (4.4)	17.8 (4.2)	-15.8**
Bond	15.9 (1.5)	15.3 (2.1)	16.1 (2.0)	-0.8
Equity Premium	-3.1 (2.8)	-13.3** (4.5)	1.7 (3.4)	

Table VIII**Average Consumption Across States and Over Odd/Even Periods.**

This table presents average consumption across states and odd/even periods. Consumption is defined as end-of-period cash holdings, in dollars. Results are stratified by participant type (autarky numbers in parentheses). Last two rows: p levels of the contribution of state (High/Low) and period (Odd/Even) to explaining variation of the consumption share of Type I (end-of-period cash holdings as a proportion of total cash available) in a two-way mixed-effects ANOVA. For choices to be Pareto efficient, consumption shares should be independent of state and period, provided the representative agents for the two types have homothetic utility.

	States		Periods	
	High	Low	Odd	Even
Type I	14.93 (19.75)	7.64 (4.69)	7.69 (2.41)	13.91 (20.65)
Type II	15.07 (10.25)	12.36 (15.31)	14.72 (20)	11.74 (5)
ANOVA p -value	0.09		0.27	
ANOVA Interaction p -value	0.23			

Table IX**End-of-Period Asset Holdings Of Type I Subjects**

This table presents average end-of-period asset holdings of Type I subjects across all replications. Initial allocation are in parentheses, for reference.

Period	Dollar Income	Asset		Total (10)
		Tree (10)	Bond (0)	
1	0	6.67	0	6.67
2	15	7.00	1.08	8.08
3	0	5.67	0.33	6.00
4	15	6.33	1.25	7.58
5	0	5.75	0.50	6.25
6	15	6.75	1.60	8.35
7	0	5.92	0.92	6.84
8	15	6.67	2.58	9.25
9	0	6.92	2.25	9.17

Table X**End-Of-Period Asset Holdings Of Three Typical Type I Subjects**

This table presents the end-of-period holdings of three typical type I subjects. Initial allocations for type I subjects are 10 Trees, 0 Bonds. Data are taken from one replication in the first Caltech session.

Asset	Subject	Period					
		1	2	3	4	5	6
Trees:							
	3	4	4	3	4	3	4
	5	1	1	0	1	1	3
	7	7	10	13	15	19	20
Bonds:							
	3	3	5	3	5	3	4
	5	8	15	14	15	16	17
	7	2	3	0	4	0	4

Table XI**GMM Estimation And Testing Results For Three Different Sets Of Instruments**

This table presents the GMM estimation and testing results for three different sets of instruments.

Instruments	χ^2 test (p value)	β (p value for $\beta = 5/6$)	γ (p value for $\gamma = 0$)
constant 1, lagged consumption growth, and asset returns	7.124 (0.310)	0.86 (0.003)	-0.01 (0.917)
constant 1, lagged consumption growth	0.731 (0.694)	0.86 (0.029)	-0.18 (0.162)
high & low state dummies, lagged consumption growth	14.349 (0.006)	0.86 (0.002)	0.16 (0.001)

Table BI**Prices, Discounts, and Equity Premia for Various Levels of CRRA.**

This table presents equilibrium prices, differences between equilibrium Bond and Tree prices (discounts), and equity premia (differences in expected returns on Tree and Bond).

CRRA (γ)	State	Tree		Bond		Price Discount	Equity Premium
		Price	Return	Price	Return		
0.2	High (H)	\$2.50	16.1%	\$2.61	15.3%	\$0.11	0.8%
	Low (L)	\$2.31	25.9%	\$2.40	25%	\$0.09	0.9%
0.5	High (H)	\$2.50	10.8%	\$2.78	8.8%	\$0.28	2%
	Low (L)	\$2.04	35.7%	\$2.27	33.3%	\$0.23	2.5%
1	High (H)	\$2.50	3.3%	\$3.13	-0.7%	\$0.63	4%
	Low (L)	\$1.67	55%	\$2.08	49%	\$0.41	6%

Table BII**Type I Agent Equilibrium Holdings and Trades.**

This table presents equilibrium holdings and trades for Type I agents. Holdings and trades are shown as a function of period (Odd/Even) and CRRA (γ). Type I agents receive income in even periods only. Calculations assume that the state in period 1 is high.

γ	Period	Tree	Bond	(Total)
0.2	Odd	5.45	2.70	(8.15)
	Even	4.63	6.23	(10.86)
	(Trade in Odd)	(+0.82)	(-3.53)	(-2.71)
0.5	Odd	6.32	1.96	(8.28)
	Even	3.48	7.24	(10.72)
	(Trade in Odd)	(+2.84)	(-5.28)	(-2.44)
1	Odd	7.57	0.62	(8.19)
	Even	2.03	7.78	(9.81)
	(Trade in Odd)	(+5.54)	(-7.16)	(-1.62)

Figures

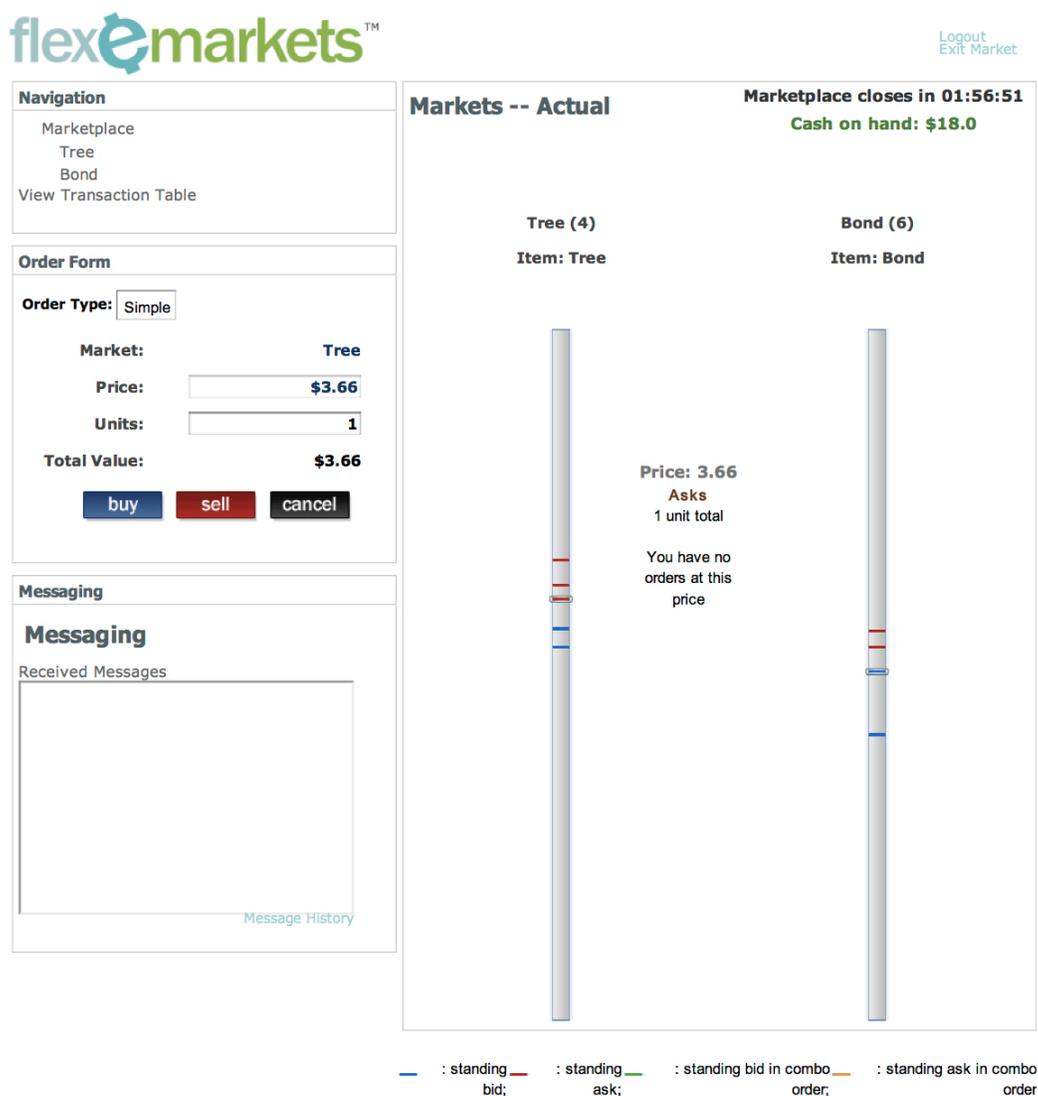


Figure 1. Snapshot of the trading interface. Two bars graphically represent the book of the market in Trees (left) and in Bonds (right). Black tags indicate standing asks; grey tags indicate standing bids. Detailed information about standing orders is provided by clicking along either of the bars (here, the Tree bar is clicked, at a price level of \$3.66). At the same time, this populates the order form to the left, through which subjects could submit or cancel orders. Asset holdings are indicated next to the name of the market, and cash balances are given in the top right corner of the interface. The remaining functionality in the trading interface is useful but nonessential. Actual interface is in color. Reproduced here in grayscale.

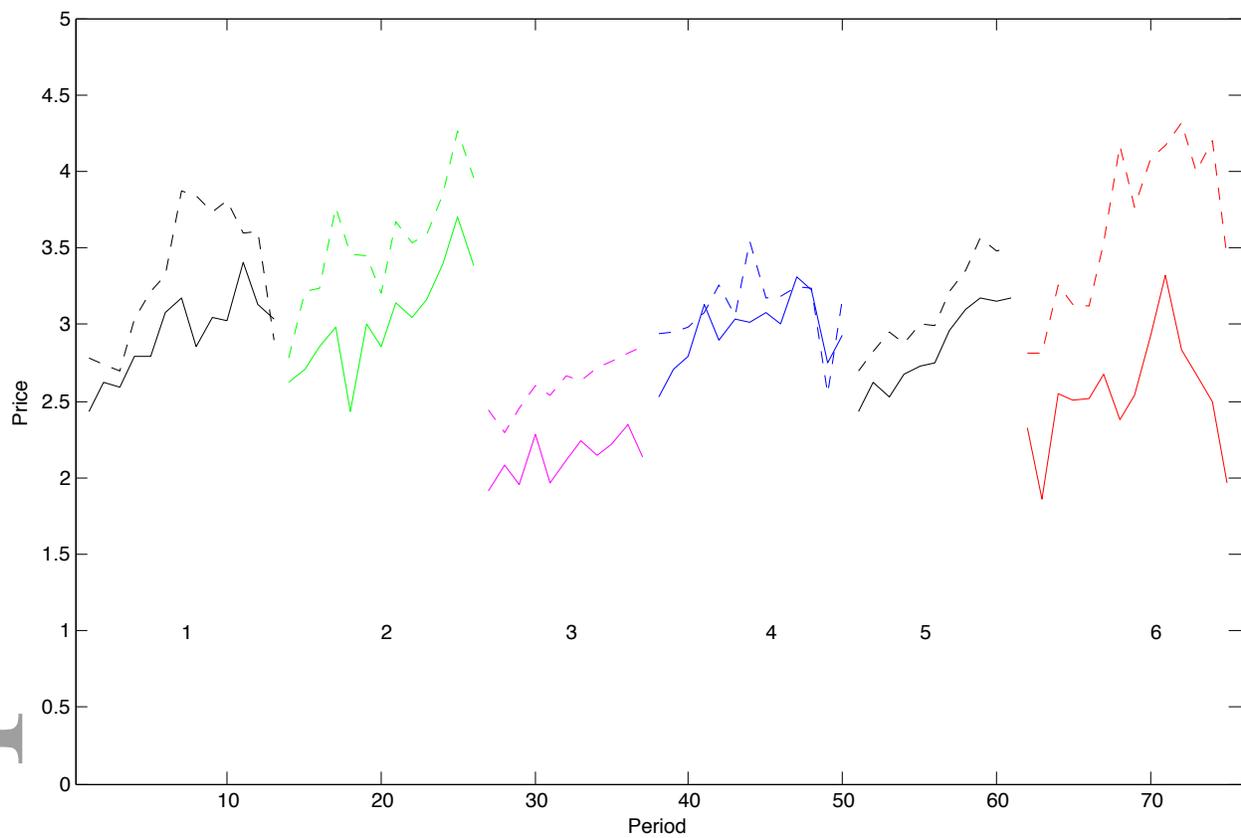


Figure 2. Time series of transaction prices. Solid line: Tree; dashed line: Bond. Averages per period. Session numbers underneath line segments refer to Table I.

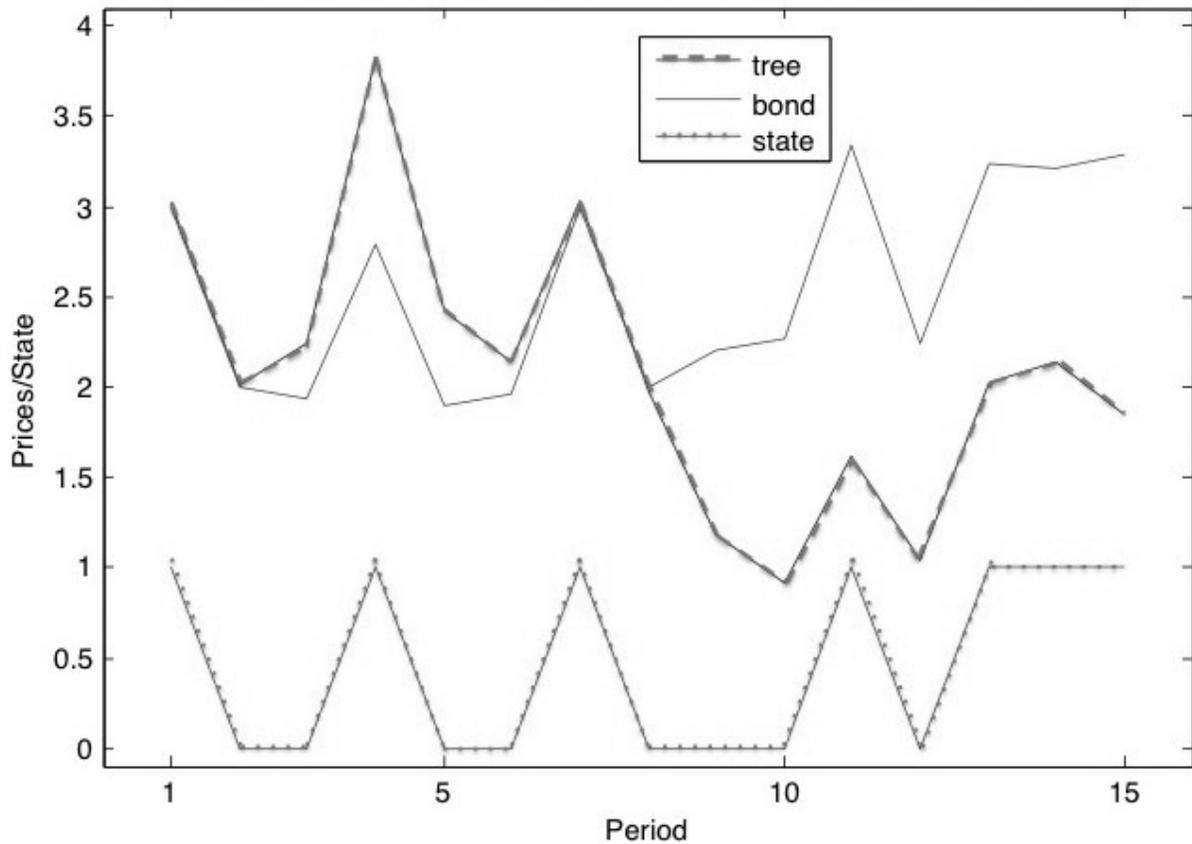


Figure 3. Time series of Tree and Bond prices in a temporary equilibrium. Solid line: Tree prices; dashed line: Bond prices. Agents expect prices to revert back to last period's levels, plus mean-zero gaussian noise with \$0.40 standard deviation. Also shown, dotted line, is the evolution of the state (High = 1; Low = 0).

Appendix: Time Line Plot To Complement Instructions

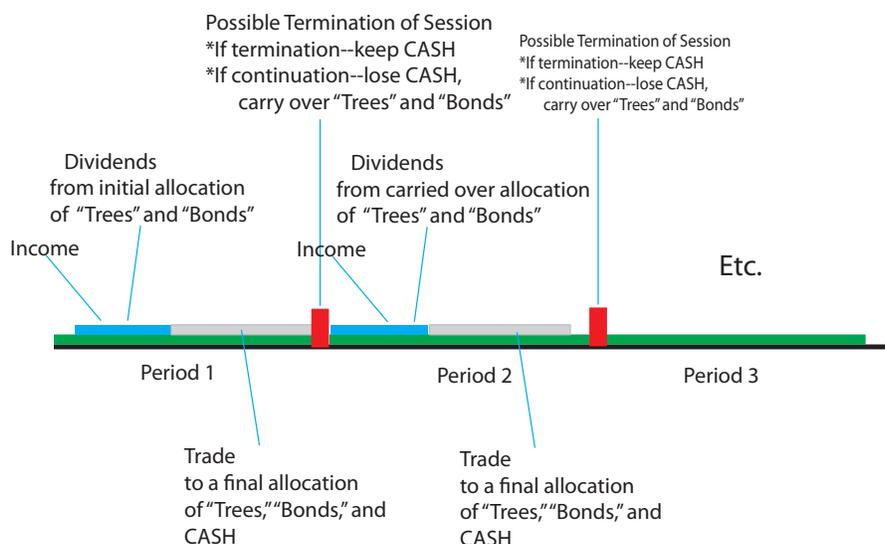


Figure C1. Timeline of Events during One Replication.
References

Klaus Adam, Albert Marcet, and Juan Pablo Nicolini. Stock market volatility and learning. Working paper, 2012.

Elena Asparouhova, Peter Bossaerts, and Charles Plott. Excess demand and equilibration in multi-security financial markets: The empirical evidence. *Journal of Financial Markets*, 6:1–21, 2003.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004. ISSN 1540-6261. doi: 10.1111/j.1540-6261.2004.00670.x. URL <http://dx.doi.org/10.1111/j.1540-6261.2004.00670.x>.

Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53, 2001. ISSN 00335533. URL <http://www.jstor.org/stable/2690442>.

Table 1.

PERIOD	1		2		3		4		5		6	
State	H		L		L		H		L		H	
Initial Holdings												
Tree	10		10		10		10		10		10	
Bond	0		0		0		0		0		0	
Dividends												
Tree	$\$1 \cdot 10 = 10$		$\$0 \cdot 10 = 0$		$\$0 \cdot 10 = 0$		$\$1 \cdot 10 = 10$		$\$0 \cdot 10 = 0$		$\$1 \cdot 10 = 10$	
Bond	$\$0.5 \cdot 0 = 0$											
Income	0		15		0		15		0		15	
Initial Cash	\$10 (=10+0+0)		\$15 (=0+0+15)		\$0 (=0+0+0)		\$25 (=10+0+15)		\$0 (=0+0+0)		\$25 (=10+0+15)	
Trade												
Tree	0		0		0		0		0		0	
Bond	0		0		0		0		0		0	
Cash Change	\$0		\$0		\$0		\$0		\$0		\$0	
Final Holdings												
Tree	10		10		10		10		10		10	
Bond	0		0		0		0		0		0	
CASH	\$ 10.00		\$ 15.00		\$ 0.00		\$ 25.00		\$ 0.00		\$ 25.00	

Table 2.

PERIOD	1		2		3		4		5		6	
State	H		L		L		H		L		H	
Initial Holdings												
Tree	10		5		6		4		5		3	
Bond	0		5		6		4		6		4	
Dividends												
Tree	$\$1 \cdot 10 = 10$		$\$0 \cdot 5 = 0$		$\$0 \cdot 6 = 0$		$\$1 \cdot 4 = 4$		$\$0 \cdot 5 = 0$		$\$1 \cdot 3 = 3$	
Bond	$\$0.5 \cdot 0 = 0$		$\$0.5 \cdot 5 = 2.5$		$\$0.5 \cdot 6 = 3$		$\$0.5 \cdot 4 = 2$		$\$0.5 \cdot 6 = 3$		$\$0.5 \cdot 4 = 2$	
Income	\$0		\$15		\$0		\$15		\$0		\$15	
Initial Cash	\$10 (=10+0+0)		\$17.5 (=0+2.5+15)		\$3 (=0+3+0)		\$21 (=4+2+15)		\$3 (=0+3+0)		\$20 (=3+2+15)	
Trade												
Tree	-5		+1		-2		+1		-2		+1	
Bond	+5		+1		-2		+2		-2		+1	
Cash Change	\$0		-\$5		+\$10		-\$7.5		+\$10		-\$5	
Final Holdings												
Tree	5		6		4		5		3		4	
Bond	5		6		4		6		4		5	
CASH	\$ 10.00		\$ 12.50		\$ 13.00		\$ 13.50		\$ 13.00		\$ 15.00	

Figure C2. Initial Holdings and Final Outcomes When Participant Does Not Trade (Top Panel) or Does Trade (Bottom Panel). Trades are assumed to take place at \$2.5 per unit. Drawing of states is for illustration only; actual drawings during the experiment may be different.

Notes

*Asparouhova is from The University of Utah, Bossaerts is from The University of Utah and The University of Melbourne, Nilanjan Roy is from The City University of Hong Kong, and William Zame is from The University of California, Los Angeles. Financial support from Inquire Europe (2006-8), the Hacker Chair at the California Institute of Technology (Caltech), the National Science Foundation (Asparouhova: SES-1061844, Bossaerts: SES-1061824; Zame SES-0079374), a Moore Foundation grant to Caltech (2006-13) in support of Experimentation with Large, Diverse and Interconnected Socio-Economic Systems, the Development Fund of the David Eccles School of Business at the University of Utah, and the Einaudi Institute for Economics and Finance is gratefully acknowledged. The paper benefited from discussions during presentations at many academic institutions and conferences. Comments from Klaus Adam, Robert Bloomfield, Luis Braido, Darrell Duffie, John Duffy, Burton Hollifield, Hanno Lustig, Ramon Marimon, Richard Roll, Jose Scheinkman, Jason Shachat, Shyam Sunder, Stijn van Nieuwerburgh, Michael Woodford, Ken Singleton (Editor), the Associate Editor and reviewers were particularly helpful.

¹This is in keeping with the prediction of static models, such as CAPM, that only market risk is priced.

² When agents exhibit homothetic utility, in equilibrium everyone consumes a constant fraction of available consumption goods.

³Crockett and Duffy (2013) test a version of the Lucas model without risk. Below we elaborate on the similarities and differences between their experimental design and the design used in this paper.

⁴The seminal paper is Hansen and Singleton (1983).

⁵The seminal paper is Mehra and Prescott (1985).

⁶See, among others, Cochrane (2001).

⁷Dunn and Singleton (1986) test a version of the model in which consumption goods provide services for two periods only.

⁸In the field, one can argue that past empirical investigations may have missed important risk components. See, for example, Heaton and Lucas (2000).

⁹Note that given the structure of the Lucas economy, residual risk can only come through the agents' imperfect price forecasts.

¹⁰Recent papers attempt to incorporate in the Lucas model uncertainty about the true nature of the evolution of fundamentals (Maenhout (2004), Epstein and Wang(1994)). However, the resulting models continue to assume perfect foresight (of prices given future states). As such, these models still do not accommodate residual price forecasting risk.

¹¹The only fundamental difference between the two settings is that in the case of differential information, agents have to form expectations about *exogenous* uncertainty, while in our setting, agents have to form expectations about *endogenous* uncertainty, namely, prices.

In the perfect foresight equilibrium, however, the problem is solved in the same way: by assuming that agents know the (endogenous) mapping from states to prices.

¹²As is well known, a stochastic ending time is (theoretically) equivalent to discounting over an infinite time horizon, assuming that subjects are expected utility maximizers with time-separable preferences.

¹³One complication that is often raised as an issue for testing models like Lucas' in the laboratory is that they require participants to be risk-averse. Consistent with evidence on individual decision making (Holt and Laury (2002)), prior experimental analysis of asset pricing models demonstrates overwhelmingly that prices reflect risk aversion as theorists know it – decreasing marginal utility. See, for example, Bossaerts, Plott and Zame (2007a) and Bossaerts and Zame (2008). This constitutes an anomaly only if one insists that participants evaluate all losses and gains in relation to the present value of lifetime wealth. Rabin (2000) cautions us about the use of a single utility function to represent preferences over all ranges of wealth.

¹⁴As Judd, Kubler, and Schmedders (2003) show, if individual endowments are stationary, then in equilibrium all trading takes place in the initial period. As Crockett and Duffy (2013) confirm, not giving subjects a reason to trade in every period (or at least frequently) is a recipe for producing price bubbles in the laboratory – perhaps because subjects are motivated to trade solely out of boredom.

¹⁵Lucas (1978) assumes that a tree and a *one-period bond* are available; we use a consol bond simply for experimental convenience. The two formulations generate the same predictions with respect to consumption and equivalent predictions with respect to prices.

¹⁶Lucas (1978) finesses the problem in a different way by defining equilibrium to consist of prices, choices, and a value function – but if unbounded debt were permitted, no value function could possibly exist.

¹⁷ This insight forms the basis of recent GMM tests on field data, whereby parameters obtained from the cross-section of asset returns are used to verify intertemporal equilibrium restrictions. See Nagel and Singleton (2011).

¹⁸To obtain the result, write all variables in terms of μ :

$$\begin{aligned} p_B^H - p_T^H &= (0.5)^2 \left(\frac{\beta}{1-\beta} \right) (\mu - 1) \\ p_B^L - p_T^L &= -(0.5)^2 \left(\frac{\beta}{1-\beta} \right) \left(\frac{1}{\mu} \right) + \text{constant} \\ p_B^H - p_B^L &= \left(\frac{\beta}{1-\beta} \right) \left(\frac{\mu}{4} \right) + \text{constant} \\ p_T^H - p_T^L &= -0.5 \left(\frac{\beta}{1-\beta} \right) \left(\frac{1}{\mu} \right) + \text{constant.} \end{aligned}$$

All variables increase in μ (for $\mu > 1$). As μ changes from one agent cohort (economy) to another, these variables all change in the same direction. Hence, across agent cohorts, they are positively correlated.

¹⁹Of course, prices *do* form a martingale under the risk-neutral probabilities – the probabilities adjusted by marginal rates of substitution – but the risk-neutral probabilities are equilibrium constructions because marginal rates of substitution depend on equilibrium allocations.

²⁰Mehra and Prescott (1985) use a slightly different model, with a long-lived tree and a one-period bond, and define the equity premium as the difference between the expected return on the risky security and that on the one-period bond.

²¹Starting with Epstein and Zin (1991), it has become standard in research on the Lucas model with historical field data to use time-nonseparable preferences, to allow risk aversion and intertemporal consumption smoothing to affect pricing differentially. Because of our experimental design, we cannot appeal to time-nonseparable preferences if we need to explain pricing anomalies. Indeed, separability over time and across states is a natural consequence of expected utility. We consider this to be a strength of our experiment: we have tighter control over preferences.

²² Our early-termination protocol relies on the specific setting we are in. Specifically, pricing in our two-period early-termination protocol is identical to that in a perpetual continuation of the random termination protocol because we are assuming that state drawing is independent over time. If we were to assume dependence, we could alter the early termination protocol. For instance, if the state transition is Markovian (meaning that the current state alone determines the transition probabilities), then a three-period early-termination

protocol could be conceived whereby the session is ended in the first, second or third stage. The key is to be able to rewrite pricing in this three-period world as an infinite series, and to show that this infinite series is identical to the pricing formula one obtains under perpetual continuation of the random termination protocol.

²³As we note earlier, we agree with Crockett and Duffy (2013) that frequent trading deters the formation of pricing anomalies such as bubbles.

²⁴Flex-E-Markets is documented at <http://www.flexemarkets.com>; the software is freely available to academics.

²⁵If a session ended with fewer replications, we paid for multiples of some or all of the replications.

²⁶In principle, subjects should be able to trade to steady-state consumption-smoothing allocations within one period; we do allow for an extra period of adjustment.

²⁷This means that the average prices of the first period of the next replication are plotted as if it were the continuation of the previous replication. The second session, for instance, has 13 observations; the first four observations pertain to the first replication, and the subsequent nine observations pertain to the second replication.

²⁸The table below displays the original estimation results. Notice that the Durbin-Watson (DW) test generated values that correspond to $p < 0.001$, suggesting that the error term was highly autocorrelated, and hence that the model was misspecified.

Original Estimation Results

OLS estimates of the regression of period-average transaction price levels on several explanatory variables, including a state dummy (* = significant at $p = 0.05$; DW = Durbin-Watson statistic of time dependence of the error term).

Explanatory	Tree Price		Bond Price	
Variables	Estim.	(95% Conf. Int.)	Estim.	(95% Conf. Int.)
Session Dummies:				
1	2.69*	(2.53, 2.84)	3.17*	(2.93, 3.41)
2	2.69*	(2.51, 2.87)	3.31*	(3.04, 3.59)
3	1.91*	(1.75, 2.08)	2.49*	(2.23, 2.74)
4	2.67*	(2.50, 2.84)	2.92*	(2.66, 3.18)
5	2.47*	(2.27, 2.67)	2.86*	(2.56, 3.17)
6	2.23*	(2.05, 2.40)	3.42*	(3.16, 3.69)
Period Number	0.06*	(0.03, 0.08)	0.06*	(0.01, 0.10)
State Dummy (High=1)	0.24*	(0.12, 0.35)	0.11	(-0.07, 0.29)
Initiate Termination	-0.07	(-0.28, 0.14)	-0.01	(-0.33, 0.31)
Dummy Even Periods	-0.00	(-0.11, 0.11)	-0.11	(-0.28, 0.06)
<hr/>				
R^2	0.71		0.52	
DW	1.05*		0.88*	

Notice, however, that the coefficient on the termination dummy is insignificant, suggesting that our termination protocol was indeed neutral, as predicted by the theory. This evidence suggests that our experimental design is correct.

²⁹We deleted observations that straddled two replications. Hence, the results in Table VI are based solely on intrareplication price behavior. Unfortunately, this deprived us from testing whether our session termination protocol was neutral. The original specification did allow us to perform this test, and as mentioned in the footnote 29, the results confirmed such neutrality.

³⁰We relate price changes to state changes using a linear model; however, because there are only two states, linearity is without loss of generality.

³¹To compute these averages, we ignored periods 1 and 2, to allow subjects time to trade from their initial holdings to steady state positions.

³²Autarky consumption of Type II subjects is independent of the state, because Type II subjects are endowed with Bonds whose dividends are riskless; autarky consumption of Type I subjects depends on states because they are endowed with Trees, whose dividends are risky. We used the sequence of realized states across all the sessions to compute their autarky consumption.

³³We implemented GMM using Matlab routines provided by Michael Cliff.

³⁴Notice the contrast between price levels, which overall are in accordance with the Lucas model, and returns, which are not. As an anonymous associate editor pointed out, price

levels may provide a more accurate reflection of subjects' expectations than (sample) average returns. If we had used, for instance, the Gordon pricing model to extract from prices the implied required rates of returns, and hence the ex ante risk premia, the evidence would have been more in line with the model.

³⁵We should add that in our experiment there is a great deal more aggregate consumption risk than in the field. In our experiment the ratio of aggregate consumption in the high dividend state to aggregate consumption in the low dividend state is 1.50, while in the field it is only about 1.08 (using Mehra and Prescott (1985) estimates of U.S. data).

³⁶Mehra and Prescott (1985) use a slightly different model, with a long-lived tree and a one-period bond, and define the equity premium as the difference between the expected return on the risky security and that on the one-period bond.

³⁷To obtain the result, write all variables in terms of μ :

$$\begin{aligned}
 p_B^H - p_T^H &= (0.5)^2 \left(\frac{\beta}{1-\beta} \right) (\mu - 1) \\
 p_B^L - p_T^L &= -(0.5)^2 \left(\frac{\beta}{1-\beta} \right) \left(\frac{1}{\mu} \right) + \text{constant} \\
 p_B^H - p_B^L &= \left(\frac{\beta}{1-\beta} \right) \left(\frac{\mu}{4} \right) + \text{constant} \\
 p_T^H - p_T^L &= -0.5 \left(\frac{\beta}{1-\beta} \right) \left(\frac{1}{\mu} \right) + \text{constant}.
 \end{aligned}$$

All variables increase in μ (for $\mu > 1$). As μ changes from one agent cohort (economy) to another, these variables all change in the same direction. Hence, across agent cohorts, they are positively correlated.

³⁸From equation A1, one can derive the (shadow) price of a one-period pure discount bond with principal of \$1, and from this price, the one-period risk-free rate. (For instance, if risk aversion is equal to one (logarithmic utility), then in the high state the one-period risk-free rate is -4% and in the low state it is 44%.) The risk-free rate mirrors changes in expected returns on the Tree and Bond. The reader can easily verify that, when defined as the difference between the expected return on the market portfolio (the per-capita average portfolio of Trees and Bonds) and the risk-free rate, the equity premium is countercyclical, just like it is when defined as the difference between the expected returns on the Tree and on the Bond.

³⁹Notice that equilibrium holdings and trade depend on whether the period is odd/even but *not* on the state (dividend of the Tree).

⁴⁰In this table, we have chosen the state in period 1 to be H so that the Tree pays a dividend of \$1. If the state in period 1 were L , and risk aversion were strictly greater than 0.5, agents would need to short sell Bonds, which we do not permit in the experiment.