

‘Lucas’ In The Laboratory

NBER Working Paper 19068^{*}

Elena Asparouhova[†] Peter Bossaerts[‡] Nilanjan Roy[§]
William Zame[¶]

May 22, 2013

Abstract

This paper reports on experimental tests of an instantiation of the Lucas asset pricing model with heterogeneous agents and time-varying private income streams. Central features of the model (infinite horizon, perishability of consumption, stationarity) present difficult challenges and require a novel experimental design. The experimental evidence provides broad support for the qualitative pricing and consumption predictions of the model (prices move with fundamentals, agents smooth consumption) but sharp differences from the quantitative predictions emerge (asset prices display excess volatility, agents do not hedge price risk). Generalized Method of Moments (GMM) tests of the stochastic Euler equations yield very different conclusions depending on the instruments chosen. It is suggested that the qualitative agreement with and quantitative deviation from theoretical predictions arise from agents’ expectations about future prices, which are *almost* self-fulfilling and yet very different from what they would need to be if they were *exactly self-fulfilling* (as the Lucas model requires).

^{*}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 Sheinkman, Jason Shachat, Shyam Sunder, Stijn van Nieuwerburgh and Michael Woodford were particularly helpful.

[†]University of Utah

[‡]Caltech

[§]Caltech

[¶]UCLA

1 Introduction

For over thirty years, the Lucas 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 particular, the price of an asset need not follow a martingale (with respect to the true probabilities) and need not be the discounted present value of its expected future dividends (with respect to the true probabilities).²

The most familiar version of the Lucas model assumes a representative agent, whose holdings consist of the aggregate endowment of securities and whose consumption is the aggregate flow of the (perishable) dividends. The representative agent has rational expectations, and so correctly forecasts both future prices and his own future decisions. Asset prices are constructed as shadow prices with respect to which the representative agent would have no incentive to trade. The heterogeneous agent version of the Lucas model that we study here assumes that all agents have rational expectations, and so correctly forecast both future prices and their own future decisions. Asset prices and allocations (consumption choices) are constructed in equilibrium. The representative agent version of the model and the heterogeneous model make the same (qualitative) price predictions; the multi-agent model also makes allocational predictions (consumption smoothing, Pareto optimality).

This paper reports on experimental laboratory tests of the Lucas model with heterogeneous agents. Our experiments provide broad support for the qualitative pricing and allocational predictions: prices are correlated with fundamentals, agents smooth consumption and insure against dividend risk. However our experiments also find that asset prices are significantly more volatile than can be accounted for by fundamentals (fundamentals explain only a small fraction of the variation of price changes of the risky asset – the Tree – and we cannot reject the null that price changes in the riskless asset – the Bond – are entirely random and unrelated to fundamentals), and agents do not insure against price risk. The data suggest that the divergence from theoretical predictions from subjects’ forecasts about future asset prices, which appear

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

²These predictions are especially important because they contradict the strictest interpretation of the Efficient Markets Hypothesis (Samuelson, 1973; Malkiel, 1999; Fama, 1991). Note that, because prices do not admit arbitrage, the Fundamental Theorem of Asset Pricing implies the existence of *some* probability measure with respect to which prices *do* follow a martingale – but that is a tautology, not a prediction.

to be vastly at odds with the predictions of the Lucas model, yet *almost self-fulfilling*. Of course asset price forecasts that are *exactly* self-fulfilling must necessarily coincide with the prices predicted by the Lucas model: this is just the definition of equilibrium. Surprisingly, however, asset price forecasts can be almost self-fulfilling and yet far from the predictions of the Lucas model – and far from equilibrium prices. This suggests that excess volatility of asset prices might not be troubling if the object of concern is welfare.

Up to now, analysis of the Lucas model – both empirical and theoretical – has focused on the “stochastic Euler equations” that deliver the equilibrium pricing restrictions (Cochrane, 2001). These equations derive from the first-order conditions of the consumption/investment optimization problem of the representative agent in the economy. It seems fair to say that empirical tests of the stochastic Euler equations using historical field data have been disappointing; indeed, beginning with Mehra and Prescott (1985), the fit of model to data has generally been considered to be poor. Attempts to improve the fit of the model to data have taken many forms. Some of this work has focused on the preferences of the representative agent, positing time-inseparability (Epstein and Zin, 1991) or loss aversion (Barberis, Huang, and Santos, 2001) or disappointment aversion (Routledge and Zin, 2011). Some of this work has focused on the nature of the data, offering corrections to the assumptions about the consumption process (Hansen and Singleton, 1983), emphasizing the role of durable goods (Dunn and Singleton, 1986) or the the role of certain goods as providing collateral as well as consumption (Lustig and Nieuwerburgh, 2005). And some of this work has focused on the statistical properties of the consumption process (Bansal and Yaron, 2004).

By contrast, our experimental study of the Lucas model focuses, not on any of these issues, but on the primitives of the model itself. In the laboratory, we can examine all predictions of the model – both the consumption predictions and the price predictions – and we are not limited to examining whether prices satisfy some set of stochastic Euler equations. This is possible because the laboratory environment allows us to observe (or control) structural information that is impossible to glean from historical data, such as the true dividend and consumption processes, agents’ beliefs about these processes, and private income flows.³

The nature of the Lucas model presents a number of unusual challenges for the laboratory environment. The most obvious challenge is that the familiar version assumes a representative agent, presumably as a shortcut to a tractable model rather than as an

³Note the similarity to the Roll (1977) critique.

assumed feature of reality. However, unless agents are identical, which seems hardly more likely in the laboratory than in the world, the representative agent is only an equilibrium construct, and not a testable assumption/prediction. Fortunately for us, the heterogeneous agent version of the 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 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. For the market outcome to be a Pareto optimal would seem to require that the market reach a Walrasian equilibrium, which in turn would seem to require a complete set of markets, an impossibility in an infinite-horizon economy with uncertainty. However, it is in fact 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), that participants are able to properly forecast future prices (as is required in a Radner perfect foresight equilibrium) and that agents can employ investment strategies that exhibit the *hedging* features that are at the core of the modern theory of derivatives analysis (Black and Scholes, 1973; Merton, 1973a) and dynamic asset pricing (Merton, 1973b).

The second challenge is that agents must learn a great deal. However, in contrast to the literature on “learning rational expectations equilibrium,” agents in our experimental economy do not need to learn/forecast the *exogenous* uncertainty – the dividend process; it is told to them. However they still must learn/forecast the *endogenous* uncertainty – the price process. As we shall see and discuss, this presents agents with a very difficult problem indeed.

In addition to these, three other particularly challenging aspects of the Lucas model need to be addressed before one can test it in the laboratory. The model assumes that the time horizon is infinite and that agents discount the future, that agents prefer to smooth consumption over time, and that 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.⁴ We provide an incentive for participants to

⁴As is well-known, a stochastic ending time is (theoretically) equivalent to discounting over an infinite time horizon (assuming subjects are expected utility maximizers with time-separable preferences).

smooth consumption by emulating perishability of consumption in each period: at the end of every non-terminal 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 accounts (and hence "consumed"). Stationarity of the economy might seem to present no difficulty and stationarity of the dividend process does indeed present no difficulty – but stationarity of the termination probability presents a very severe difficulty. If an experimental session lasts for two hours and each period within that session lasts for four minutes, it is 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 is quite hard for participants to believe that the termination probability remains the same at the end of the twenty-ninth period, when 116 minutes have elapsed and only four minutes remain. In that circumstance, participants will surely believe – correctly – that the termination probability must be higher at the end of the session. However, if subjects believed the termination probability is not constant, a random ending time would induce a non-constant discount factor – and very likely induce different discount factors in different subjects. To treat this challenge, we introduce a novel termination rule.

As is the case in most (all?) asset-pricing models, behavior in the Lucas model is driven by risk aversion. Because laboratory stakes are small – it would be rare for a subject to earn as much as \$100 – a natural concern is whether risk aversion is observable in the laboratory. If, as is often assumed in the literature, subjects evaluate all losses and gains in relation to present value of lifetime wealth – and keeping in mind that \$100 is surely less than 0.1% of present value of lifetime wealth for virtually all subjects, and much less for most – it would seem that risk aversion could not be observable in the laboratory. However, there is ample evidence that subjects *do* display substantial risk aversion in the laboratory (presumably because they *do not* evaluate all losses and gains in relation to present value of lifetime wealth); see Holt and Laury (2002) for risk aversion in laboratory betting environments, Bossaerts and Zame (2008) for risk aversion in laboratory asset markets, and Rabin (2000) for cautions about the use of a single utility function to represent preferences over all ranges of wealth.

In parallel work, Crockett and Duffy (2013) also study an infinite horizon asset market in the laboratory, but their experimental approach and purpose are very different from ours. First, there is *no risk* in their setting. Our experiment includes risk, and hence, allows us to study the interplay between risk avoidance and inter-temporal smoothing. This interplay generates the stochastic Euler equations that informed the

empirical research on the Lucas model and from which have emerged the controversies that we alluded to before, about the historical magnitude of the equity premium, historical asset price volatility, and the presence of return predictability in historical data. As noted, Crockett and Duffy (2013) study only consumption smoothing. Second, Crockett and Duffy (2013) *induce* a preference for consumption smoothing by imposing a schedule of final payments to participants that is non-linear in period earnings. A problem with that approach – aside from the question of whether one should try to induce preferences rather than take them as given – is that this is (theoretically) equivalent to time-separable additive utility *only* if participants’ true preferences are risk-neutral – but there is ample laboratory evidence that participants display substantial risk-aversion even for relatively small laboratory stakes, as mentioned before. Third, the focus in Crockett and Duffy (2013) is on *incentives to trade*, contrasting pricing in a treatment where there are incentives because of demand for consumption smoothing against a treatment where there are none because there is no demand for consumption smoothing. As such, Crockett and Duffy (2013) use their treatment to shed light on a long line of experimental work on asset price bubbles. Indeed, their treatment without incentives is identical to the one that is known to trigger bubbles, as first documented in Smith, Suchanek, and Williams (1988). In contrast, here we are interested in aspects of the Lucas model that have generated controversy in studies of historical asset prices. For good experimental control, our setting is one where there are definite incentives to trade. A final difference concerns the *stationarity* assumption in the Lucas model. Both Crockett and Duffy (2013) and we use random termination to induce discounting, but even with constant termination probability, one inevitably must confront the fact that the experiment must end at a pre-announced point in time, as explained before. This issue is ignored in Crockett and Duffy (2013); in contrast, we recover stationarity through a novel experiment termination protocol.

Some of our colleagues have wondered why anyone would bother to carry out laboratory tests of an asset pricing model that is rejected on the basis of historical field data. To us, a more natural question would seem to be why anyone would bother to carry out laboratory tests of an asset pricing model that is accepted on the basis of historical field data (not to mention that the latter set of models would seem small – or empty). Models represent ideal environments and should not be expected to fit perfectly to a non-ideal world. The laboratory is as close as we can come to the ideal environment that models are intended to represent. Understanding the performance of a model in the laboratory – the extent of its success or failure in all the dimensions in which it is predictive/descriptive in an ideal environment – may tell us how and

whether the model is or is not predictively/descriptively useful in a non-ideal world. This seems especially true of asset pricing models that describe/predict both prices and choices because both prices and choices can be observed in the laboratory data – but choices cannot be observed in historical field data.⁵ We choose to test the basic Lucas asset pricing model – rather than variants such as those of (Mehra and Prescott, 1985) – not because it most closely resembles the world, but because it is clean and simple and because its predictions are driven by precisely the same forces that drive the predictions of more complicated variants. If those forces cannot be observed at work in the Lucas model, we cannot see why one should expect them to be observed at work in more complicated variants.⁶

The remainder of this paper is organized as follows. Section 2 presents the Lucas model within the framework of the laboratory economy we created. Section 3 provides details of the experimental setup. Results are provided in Section 4. Section 5 discusses potential causes behind the excessive volatility of asset prices observed in the laboratory markets. Section 6 examines the laboratory data through the lens of the statistical analysis that has traditionally been employed on historical field data. Section 7 concludes.

2 The Lucas Asset Pricing Model

We formulate a particular instantiation of the Lucas asset pricing model that is simple enough to implement 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 preferences and endowments (of assets) and time-varying consumption streams, and we make no assumptions about functional forms

⁵As we have noted earlier, many parameters can also be observed – or determined – in the laboratory but not in the field.

⁶Models in which the growth rates of dividends (rather than the levels) are stationary seem difficult to or impossible to test in the laboratory in part because growth is difficult to handle smoothly in the laboratory – both because of participant perceptions and because payoffs depend on the duration of a replication – and in part because their central predictions rely on assumptions that seem unlikely to hold in the laboratory environment. The Mehra and Prescott (1985) model, for instance, derives its pricing predictions from the assumption that there is a representative agent who has homothetic preferences; absent these assumptions, it is unclear what pricing predictions can be drawn. As we show in our discussion of the Lucas model in Subsection 2.2, we can derive strong and testable predictions about both prices and choices even in the absence of a representative agent or of any assumptions about the preferences of agents (beyond expected utility and risk aversion).

– but still obtain strong and testable implications for individual consumption choices and trading patterns and for prices.

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 may not be.⁷

We caution the reader (again) that our formulation assumes stationarity in the *levels* of dividends and aggregate consumption rather than in *growth rates*, as in Mehra and Prescott (1985) and much subsequent work that has used historical field data to inform empirical research. We choose stationarity in levels because it is easier to implement and easier for subjects to understand, which would seem desirable – perhaps necessary – criteria for an experiment that already poses many other challenges. As we show below, stationarity in levels also has the feature that it leads to strong and testable predictions, even in the presence of heterogeneity across subjects and the absence of any assumptions about functional forms; stationarity in growth rates does not seem to have this feature.

2.1 A General Environment

We consider an infinite horizon economy with a single consumption good in each time period. In the experiment, the consumption good is cash so we use the terms ‘consumption’ and ‘cash’ interchangeably here. In each period there are two possible states of nature H (high), L (low), which occur with probabilities $\pi, 1 - \pi$ independently 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 , 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.⁸ We as-

⁷As Judd, Kubler, and Schmedders (2003) has shown, if individual endowments are stationary then, at 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.

⁸Lucas (1978) assumes that a *Tree* and a *one-period bond* are available; we use a *consol bond* simply for experimental convenience.

sume $d_T^H > d_T^L \geq 0$ 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 dividends processes are stationary in levels. (In the experiment proper, we choose $\pi = 1/2$; $d_T^H = 1$, $d_T^L = 0$; $d_B = 0.50$, with all payoffs in dollars.)

There are n agents. 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 – i.e., a time-invariant function of dividends (in the experiment proper it will be constant) – so that aggregate consumption $b + \tau + e$ is also stationary, 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 assume that each 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 0 (so that optimal consumption choices are interior), but 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 . (In the experimental setting this seems an especially reasonable assumption because the discount factor is just the probability of continuation, which is constant and common across agents.)

In each period t agents receive dividends from the Bonds and Trees they hold, 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. Agents take as given the current prices of the Bond $p_{B,t}$ and of the Tree $p_{T,t}$ (both of which depend on the current state), make forecasts of (stochastic) future asset prices $p_{B,t'}, p_{T,t'}$ 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. 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 that equilibrium be well-defined, such schemes must be ruled out. Levine and Zame (1996), Magill and Quinzii (1994) and Hernandez and Santos (1996) show that this can be done in a number of different 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.⁹ (In the experimental setting, we forbid short sales so debt is necessarily bounded.)

As is universal in the literature we assume that a Radner equilibrium exists and – because markets are (potentially) dynamically complete – that it coincides with Walrasian equilibrium and in particular 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 all that we are assuming is *subsumed* in the familiar version.

2.2 Predictions

We first show that, despite allowing for heterogeneity and without making any assumptions about functional forms, the model makes testable *quantitative* predictions about individual consumptions, prices and trading patterns. These predictions, which are entirely familiar in the context of the usual Lucas model with a representative agent having CRRA utility, follow from the nature of uncertainty and the assumption of Pareto optimality. Some of these predictions take a particularly simple form when the specific parameters are as in the experiment. In Subsection 2.3 below we provide explicit solutions in the case that agents all possess constant relative risk aversion utility.

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 $t, t + 1$: $u'_i(c_{i,t+1}^\sigma)/u'_i(c_{i,t}^\sigma) = u'_j(c_{j,t+1}^\sigma)/u'_j(c_{j,t}^\sigma)$ for each i, j . In particular, the ranking of marginal utility for consumption in state σ at dates $t, t + 1$ must be the same for all agents. Because utility functions are strictly concave, the rankings of consumption in state σ at dates $t, t + 1$ must be the same

⁹Lucas (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 then no value function could possibly exist.

for all agents (and opposite to rankings of marginal utilities). But the sum of individual consumptions is aggregate consumption, which is stationary – hence equal in state σ at periods $t, t + 1$. Hence the consumption of each individual agent must also be equal in state σ at periods $t, t + 1$. Since t is arbitrary this means that individual consumption must be constant in state σ ; i.e., 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.

To see this, fix an agent i ; 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, 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 $A = B, T$. (1) 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 (1) is independent of the choice of agent i , and also that we could write (1) 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 (1) 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

$$\begin{aligned}
p_{A,t}^\sigma &= \sum_{\tau=0}^{\infty} \beta^{\tau+1} E \left[\frac{u'_i(c_{i,t+\tau+1})}{u'_i(c_{i,t}^\sigma)} d_{A,t+\tau+1} \right] \\
&= \beta E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right] \sum_{\tau=0}^{\infty} \beta^\tau \\
&= \left(\frac{\beta}{1-\beta} \right) E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right] \tag{2}
\end{aligned}$$

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 (2) 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

$$p_A^H/p_A^L = \mu \tag{3}$$

Substituting and solving yields

$$\begin{aligned}
p_A^H &= \left(\frac{\beta}{1-\beta} \right) \left[\pi d_A^H + (1-\pi)d_A^L \mu \right] \\
p_A^L &= \left(\frac{\beta}{1-\beta} \right) \left[\pi d_A^H (1/\mu) + (1-\pi)d_A^L \right] \tag{4}
\end{aligned}$$

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

$$p_B^H = (2.5)(1 + \mu)/2 \tag{5}$$

$$p_B^L = (2.5)(1 + \mu)/2\mu \tag{6}$$

$$p_T^H = 2.5 \tag{7}$$

$$p_T^L = 2.5/\mu \tag{8}$$

In particular, $p_T^H = 2.5$ (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 correlated with fundamentals.

This is also an immediate consequence of equations (4); 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 (4). 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. The equity premium is positive and counter-cyclical.

The difference in the prices of the Tree and the Bond can be translated into differences in expected returns; 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).¹⁰ 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 , $p_A^H = \mu p_A^L$):

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

¹⁰Mehra and Prescott (1985) use a slightly different model, with long-lived Tree and a one-period bond.

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

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 counter-cyclical (lower in the High state than in the Low state). Note that counter-cyclicity 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.

Conversely, the discount of the price of the Tree relative to that of the Bond $p_B - p_T$ is pro-cyclical. (This follows directly from the fact that the ratio of the prices across states of both securities are equal and the fact that the Bond is always more expensive than the Tree.)

8. Asset prices and returns are predictable.

Asset prices are predictable because they depend on the state; again this is embodied in (4). That returns are predictable follows from the additional fact that dividends are i.i.d. It is important to note that predictability of prices and returns flatly contradicts the simplest versions of the Efficient Markets Hypothesis, which asserts that prices form a martingale under the *true* probabilities (Samuelson, 1973; Malkiel, 1999; Fama, 1991). (Predictability of asset prices and returns was one of the original points made by Lucas (1978).) Of 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*.

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.¹¹

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 price risk.

If there were no price risk, agents could smooth consumption simply by buying or selling one asset. However, there *is* price risk, because prices move with fundamentals and fundamentals are risky. 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 counter-intuitively! – agents *buy* Trees in periods when private income is low and *sell* when private income is high.

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

It can be shown that price risk hedging increases with the risk aversion of the

¹¹To 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.

representative agent. This is because equilibrium price risk, measured as the difference in prices across H and L states, increases with risk aversion (embedded in μ).

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. Investment choices should be such that consumption (cash holdings at the end of a period) across states becomes perfectly rank-correlated between agent types (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 others, that agents hedge price risk, by buying Trees when experiencing income shortfalls (and selling Bonds to cover the shortfalls), and selling Trees in periods of high income (while buying back Bonds).

2.3 A Numerical Example

For illustration, we compute predicted equilibrium prices, holdings and consumptions, taking the parameters as in the experiment and assuming that all agents display identical constant relative risk aversion.

- There are an even number $n = 2m$ of agents; agents $i = 1, \dots, m$ are of Type I, 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 constant relative risk aversion $\gamma = .2, .5, 1$. (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 $.2 - .5$ is consistent with the experimental findings of Holt and Laury (2002) and Bossaerts and Zame (2008).
- The initial state is High.

Table 1 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.

Table 1: Prices, discounts and equity premia for various levels of constant relative risk aversion (γ).

γ	State	Tree		Bond		Price	Equity
		Price	Return	Price	Return	Discount	Premium
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%

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 pro-cyclical behavior of the discount and the counter-cyclical behavior of the equity premium. The dependence of prices on the state, and the predictability of returns is apparent from the table.¹²

Table 2 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 across time.^{13,14}

¹²From Equation 1, 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 1 (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 return on the Tree and on the Bond.

¹³Notice that equilibrium holdings and trade depend on whether the period is odd/even but *not* on the state (dividend of the Tree).

¹⁴In 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.

Table 2: Type I agent equilibrium holdings and trades as a function of period (Odd/Even) and constant relative risk aversion (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)

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 have noted earlier, without the assumption of identical homothetic utilities all we can conclude is that individual consumptions are perfectly correlated with aggregate consumption.

3 Implementing the Lucas Model

As we have already noted, implementing the Lucas economy in the laboratory encounters 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 used 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 closed we rolled a twelve-sided die; if it came up either 7 or 8, we terminated; otherwise we moved on to a new period.

To resolve issue (b), we made end-of-period individual cash holdings *disappear* in every period that was not terminal; only securities holdings carried over to the next period. If a period *was* terminal, however, securities holdings perished and cash holdings were credited; participants' earnings were then determined entirely by the cash they held 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.¹⁵

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 non-negligible 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

¹⁵Starting with Epstein and Zin (1991), it has become standard in research on the Lucas model with historical field data to use time-nonseparable preferences, in order 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 across time and states is a natural consequence of expected utility. We consider this to be a strength of our experiment: we have tighter control over preferences. This is addition to our control of beliefs: we make sure that subjects understand how dividends are generated, and how termination is determined.

approached. To deal with this problem we employed a simple termination rule: We announced that the experimental session would last until a pre-specified time and there would be as many replications of the (Lucas) economy as could be fit within this time frame. If a replication ended at least 10 minutes before the announced ending time of the session, a new replication would begin; otherwise, the experimental session would end. *If a replication was still running 10 minutes before the announced ending time of the session, we announced before trade opened that the current period would be either the last one (if our die turned up 7 or 8) or the next-to-last one (for all other values of the die). In the latter case, the next period was the terminal period, with certainty, so subjects would keep the cash they received through dividends and income for that period.* (There should be no trade in the terminal period because assets perish at the end and hence are worthless.) In the Appendix, we re-produce the time line plot that we used alongside the instructions to facilitate comprehension.

To see that equilibrium prices remain the same whether the new termination protocol is applied or if 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 expected marginal utility from keeping cash worth one unit of the security; the right-hand side is expected marginal utility from buying the unit; optimality implies equality. Re-arranging 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 dividends, consumption and prices are stationary, this reduces to (2), as asserted.

In the experiment, the task for the subjects is to trade off cash against securities. In a given period, cash is desirable because it constitutes experimental earnings if the current period is in fact the terminal period; securities are desirable because they generate experimental earnings in future periods if the current period is in fact not the terminal period. It was easy for subjects to grasp the essence of this task, and the simplicity allowed us to make instructions short. See the Appendix for sample instructions.

There is one further difficulty which we have not mentioned: default. In the (finite or infinite horizon) Radner model, assets are simply promises; selling an asset – borrow-

ing – entails a promise to repay in the future. However, in the model, nothing enforces these promises: that they are kept in equilibrium is simply part of the definition of equilibrium. If nothing enforced these promises in the laboratory then participants could (and in our experience, would) simply make promises that they could not keep. One possibility for dealing with this problem is to impose penalties for default – 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: What should the punishment 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 an 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, this speaks only to intentional default, but what about *unintentional default – mistakes*? And what about plans that would have led to default in circumstances that might have occurred but did not? And 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 (2007) (and others) have pointed 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 the analysis of Section 2 presumes that it was 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 0 he could not sell it and if his consumption and portfolio were both 0 he could not buy it. However, so long as agents do not bump up against the zero bound, the analysis remains correct; in the actual experimental data, the number of agents who bumped up against the zero bound was quite small. In our analysis, therefore, we shall simply take note of the prohibition of short sales but assume that the prohibition is never binding.

The need to bar short sales explains why we use an instantiation 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 was not permitted). Because both assets are in positive supply, our economy is, strictly speaking, a Lucas *orchard* economy

(Martin, 2011), but the qualitative predictions of the model are not different from those of a model in which the Bond is in zero net supply.

Because income and dividends vary across time and states and cash disappears at the end of each non-terminal period, those who are optimizing or nearly optimizing must trade often. As we shall see, trading volume was indeed substantial.¹⁶ 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.¹⁷ Subjects quickly familiarized themselves with key aspects of trading in the open-book mechanism (bids, asks, cancelations, 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 re-produced 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.

Within each experimental session, we conducted as many replications as possible within the time allotted. In order to minimize wealth effects, we paid for only a fixed number (say, 2) of the replications, randomly chosen after conclusion of the experiment.¹⁸ 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 give session – and they are certainly not constant across sessions, since the populations of subjects in different sessions were disjoint.

4 Results

Table 3 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. 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 and hence, generated no trade. In the other (unstarred) session, the last replication occurred sufficiently close to the end of the experiment that a new

¹⁶As we have noted earlier, we agree with Crockett and Duffy (2013) that frequent trading deters the formation of pricing anomalies such as bubbles.

¹⁷Flex-E-Markets is documented at <http://www.flexemarkets.com>; the software is freely available to academics upon request.

¹⁸If a session ended with fewer replications we paid for multiples of some or all of the replications.

Table 3: Summary data, all experimental sessions.

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)	

replication was not begun, so our termination protocol was not applied.

We first discuss volume, and then look at prices and choices.

Volume. Table 4 lists average trading volume per period (excluding terminal periods in which there should be no trade). Trading volume in Periods 1 and 2 is significantly higher, reflecting trading needed 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.¹⁹ 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 are endowed with 10 Trees and have 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 210 seconds long, this means that roughly 10% of available securities were traded each period and that one transaction occurred roughly every 5 seconds. While substantial (recall that we choose a design in which subjects *ought* to trade often), this volume of trade is low compared to theoretical predictions: Table 2 shows that even for γ (risk aversion)

¹⁹In 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.

Table 4: Trading volume.

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

as low as 0.2, average volume should be about 20% of average holdings for the Tree and over 50% for the bond. As we shall see, however, subjects did not follow the sophisticated price hedging strategy reflected in the numbers in Table 2 (which was to *buy* Trees when in need of cash). Without this hedging strategy, the volume needed to smooth consumption is substantially lower. Fewer assets have to be sold since no offset of Tree purchases is necessary (Trees are bought as a hedge for price risk).

Cross-Sectional Price Differences. Table 5 displays average period transaction prices as well as the period's state (High if the dividend of the Tree was \$1; Low if it was \$0). Consistent with the Lucas model, the Bond is priced *above* the Tree, with a price differential of about \$0.50.

Prices Over Time. Figure 2 shows a plot of the evolution of (average) prices over time, arranged chronologically by experimental sessions (numbered as in Table 3); replications within a session are concatenated. The plot reveals that prices are volatile. In theory, prices should move only because of variability in economic fundamentals,

Table 5: Period-average transaction prices and corresponding discount of the Tree price relative to the Bond price.

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

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. In reality, a large fraction of price movements is unrelated to fundamentals; following LeRoy and Porter (1981) and Shiller (1981), we will refer to this as *excessive volatility*. Some price drift can be detected, but formal tests reported below will reveal that the drift is entirely due to the impact of states on prices, and the particular sampling of the states across the sessions.

As Table 6 shows, p_T^H is higher than predicted (it is predicted to be 2.50 but is actually 2.91) and that the ratio p_T^H/p_T^L is greater than the ratio p_B^H/p_B^L ; however it is not clear whether these deviations are statistically significant. Prices in the High state are on average 0.24 (Tree) and 0.14 (Bond) above those in the Low state – prices do move with fundamentals (dividends). (We do not show statistical information because (average) transaction prices are not i.i.d., so that we cannot rely on standard t tests to determine significance. We provide formal statistical evidence later on.) Table 6 also shows that the discount on the price of the Tree relative to the Bond is higher in the Low state than in the High state; the observed discount is counter-cyclical. As we have shown earlier, the prediction of theory is that the discount should be higher in the High state than in the Low state; the predicted discount is pro-cyclical.

Cross-Sectional And Time Series Price Properties Together. The theory predicts that the differential in prices between High and Low states should correlate positively with the difference between the Bond price and the Tree price, i.e., the discount of the Tree price relative to the Bond price. Correlation is to be taken across economies, where economies are distinguished only by session cohort. Table 7 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 observa-

Table 6: Mean period-average transaction prices and corresponding discount of the Tree price relative to the Bond price, as a function of state.

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

Table 7: Correlation across replications between the average discount on the Tree price relative to the Bond price and the average price differential of the Tree or Bond between High and Low states.

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

tions in total. Consistent with the theoretical prediction, the correlations are positive, though the estimate is insignificant for the Bond.

Prices: Formal Statistics. To enable formal statistical statements about the price differences across states, we ran a regression of period transaction price levels onto the state (=1 if high; 0 if low). To adjust for time series dependence evident in Figure 2, we 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 either the penultimate or last one, depending on the draw of the die. Lastly, we add a dummy for even periods. Table 8 displays the results.

We confirm the positive effect of the state on price levels. Moving from a Low to a High state increases the price of the Tree by \$0.24, while the Bond price increases by \$0.11. The former is the same number as in Table 6; the latter is a bit lower. The price increase is significant ($p = 0.05$) for the Tree, but *not* for the Bond.

The coefficient to the termination dummy is insignificant, suggesting that our termination protocol is neutral, as predicted by the Lucas model. This constitutes comforting evidence that our experimental design was correct.

Table 8: OLS regression of period-average transaction price levels on several explanatory variables, including state dummy. (* = significant at $p = 0.05$; DW = Durbin-Watson statistic of time dependence of the error term.)

Explanatory Variables	Tree Price		Bond Price	
	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)
R^2		0.71		0.52
DW		1.05*		0.88*

Table 9: OLS regression of changes in period-average transaction prices. (* = significant at $p = 0.05$.)

Explanatory Variables	Tree Price Change		Bond Price Change	
	Estim.	(95% Conf. Int.)	Estim.	(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

However, closer inspection of the properties of the error term revealed substantial dependence over time, despite our including dummies to mitigate time series effects. Table 8 shows Durbin-Watson (DW) test statistics with value that correspond to $p < 0.001$. Therefore, inference gauged from the results displayed in Table 8 cannot be trusted.

Further model specification analysis was performed, to ensure that the error term became properly behaved. This revealed 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 the result of only the change in the state. See Table 9.²⁰ 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 the state is accounted for. This implies, among others, 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). 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 significant ($p < 0.05$) and substantial (\$0.19). The effect of a change in state on the Bond price is lower (\$0.10), though insignificant ($p > 0.05$). Both confirm the theoretical prediction that prices should be determined by the state.

The excess volatility of prices is apparent from Table 9. Fundamentals (changes in the state) explain only 18% of the variability of the Tree prices ($R^2 = 0.18$); 82% of

²⁰We deleted observations that straddled two replications. Hence, the results in Table 9 are solely based on intra-replication price behavior.

price variance is unexplained.²¹ The situation is even worse for the Bond: 96% of the variance of Bond price changes are unexplained by changes in the state. (Of course the Lucas model predicts that Fundamentals should explain *all* of the variability of asset prices.) It deserves emphasis that the unexplained variability is essentially noise; in particular, it is unrelated to the subject cohort, because session dummies were insignificant. Overall, the regression in first differences shows 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 distinct drift.

Consumption Across States. 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 10 displays the average amount of cash (consumption) per type in High vs. Low states.²² Consistent with the theoretical prediction, consumption is positively rank-correlated across Types. To gauge the significance of this finding, Table 10 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 consumption levels are *anti-correlated*. Through trading, the average Type I and Type II agents manage to move their consumptions from negatively to positively correlated, suggesting economically significant Pareto improvements, consistent with the Lucas model.

Consumption Across Odd And Even Periods. Another prediction is that subjects should be able to perfectly offset income differences across odd and even periods. Table 10 demonstrates that our subjects indeed managed 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). Therefore, the experimental results suggest substantial Pareto improvements through trading.²³

Consumption Shares Across States and Across Odd and Even Periods. If one is willing to entertain the assumption that utilities of our subjects are homothetic,

²¹We relate price changes to state changes uses a linear model; however, because there are only two states, linearity is without loss of generality.

²²To compute these averages, we ignored Periods 1 and 2, to allow subjects time to trade from their initial holdings to steady state positions.

²³Autarchic consumption of Type II subjects is independent of the state, because Type II subjects are endowed with Bonds whose dividends are riskless; autarchic 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.

Table 10: Average consumption (end-of-period cash, in dollars) across states (High or Low Tree dividend) and across periods (Odd/Even), stratified by participant Type (autarchic numbers in parentheses). Last two rows: p levels of the contribution of State and Period 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 (Analysis of Variance). For choices to be Pareto efficient, consumption shares should be independent of State and Period (provided the representative agents for the two participants 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			

Pareto efficiency suggests 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 10 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 10 shows that neither the state nor the nature of the period (nor their interaction) are significant factors ($p > 0.05$) 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 10 are solely due to sampling error.

The finding is rather striking, because the assumption of homothetic preferences is

Table 11: End-of-period asset holdings, type I subjects. Averages across all replications and subjects (of Type I). Initial allocation in parentheses, for reference.

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

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.

Price Hedging. The above results suggest that our subjects (on average) managed to move towards the Pareto-optimal equilibrium consumption patterns of the Lucas model. However, contrary to model prediction, they did not resort to price hedging as a means to ensure those patterns. Table 11 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 “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 price 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).

Subject-Level Differences There are, however, significant individual differences in portfolio choices. Table 12 illustrates how three subjects of Type I end up holding almost opposing portfolios of Trees and Bonds. Subject 7 increased his holdings of Trees over time. Significantly, this subject bought Trees even in periods with income shortfall (odd periods), effectively implementing the price hedging strategy of the theory. Subject 5 is almost a mirror image of subject 7, though s/he did not resort to price hedging. Subject 3 diversified across Trees and Bonds but did not hedge price risk either because Tree holdings decreased in odd periods.

The subject-level differences reported in Table 12 are no exception. The contrast between choices at the individual level and at the Type level is sharp. The theory “works” at the Type level, but not at the individual level. This contrast suggests that one has to be careful extrapolating to phenomena at the market level (e.g., prices) from

Table 12: End-Of-Period Asset Holdings Of Three Type I Subjects. Initial allocations: 10 Trees, 0 Bonds. Data from one replication in the first Caltech session.

Subject	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

observing individuals singly. If we had taking any of the three subjects as “typical,” 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, e.g., CAPM) even if individual choices are at odds with the theory; see Bossaerts, Plott, and Zame (2007a).

5 The Expected and the Anomalous

With respect to the predictions of the Lucas model, our experiments generate findings that are expected – prices and individual consumption are correlated with fundamentals (aggregate consumption) – and findings that seem anomalous – prices are excessively volatile and price risk is not hedged. Because volatile prices would seem to signal clearly the presence of price risk that should be hedged, the co-existence of excess volatility and un-hedged price risk 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 need or ability to hedge against price risk

To see why this might be so, recall first that the predictions of the Lucas model – and indeed the very definition of Radner equilibrium – depend on the assumption that agents have perfect foresight and in particular that the beliefs of subjects about the dividend process and the price process are *exactly correct*. In the experiment, the

subjects are told the dividend process but the price process must be learned, so it would be too much to expect that beliefs be *exactly* correct – but perhaps not too much to expect that beliefs be *approximately* correct. Optimization against exactly correct beliefs leads exactly to the Radner equilibrium predicted by the Lucas model, and 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 so: because the price process is *endogenous*, beliefs about the price process can be approximately correct and still very far from the price process predicted by the Lucas model. The same point has been made by Adam, Marcet, and Nicolini (2012), who used it to explain excess volatility in historical data (as we use it to make sense of excess volatility in our experimental 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 the (naive version of the) Efficient Markets Hypothesis – and *not* to co-move with economic fundamentals – as would be predicted by the Lucas model. This belief is wrong, but it is not readily falsifiable on the basis of the limited number of observations available to subjects. Indeed, the belief that Bond prices do not follow a martingale would not be falsifiable even after 80 observations – an order of magnitude more observations than were available to subjects. The belief that prices follow a martingale is thus a credible working hypothesis.

A thought experiment may help to understand the consequences of these incorrect beliefs. Imagine that in every period agents always believe that past prices are the best predictions of future prices, *independently 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 the market generates prices so that demand and supply are equal in that period. Of course, beliefs are wrong and will be revealed to be wrong next period, so we are considering in this thought experiment only a kind of temporary equilibrium, but one 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 do co-move with dividends, but very noisily – hence they are excessively volatile.

Figure 3 displays the evolution of prices and states in a typical simulation of this temporary equilibrium. There are two types of agents, endowed as in the experiment, each type 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. This produces price evolutions very much in line with those in the experiment. At the same time, agents do not hedge price risk (they don't perceive any and accommodate income shortfalls solely by selling Bonds *and* Trees). Still, their choices do move substantially towards Pareto optimality: the consumption share of the Type I agent fluctuates only between 39% and 44%, little affected by state and period (Odd/Even).

This thought experiment demonstrates starkly that the price predictions of the Lucas model are fragile to small mistakes in beliefs about the price process. This comes as a surprise because the price predictions of the Lucas model are robust to small mistakes in beliefs about the dividend process (Hassan and Mertens, 2010). However, it seems that mistakes in beliefs about these processes can manifest themselves quite differently: because the dividend process is exogenous, mistakes are damped out; because the price process is endogenous, mistakes can create positive feedback.

6 The Data Viewed Through A Traditional Lens

As noted before, we (control and) observe much more in our laboratory environment than is possible in the field: our laboratory data are much richer than historical field data. However, an interesting and potentially revealing exercise is to ignore this additional richness, treat our laboratory data as if it were historical field data, and carry out on our laboratory data the same kinds of econometric tests that have been carried out on historical field data. In particular, we can consider only the times series of asset returns and aggregate consumption in our laboratory data – ignoring the additional information (true dividend process, true realized state, individual choices, etc) – and use the Generalized Method of Moments (GMM) to test whether the first-order conditions (the stochastic Euler equations) are satisfied for a representative agent. To be specific, we assume – as has been done in the analysis of historical field data – the presence of a representative agent with constant relative risk aversion (power utility); we estimate the coefficient of risk aversion γ and the discount factor (β) while testing whether the Euler equations hold, assuming that the representative agent “consumes” the aggregate cash each period. Note that, because we have only two states, the assumption of constant relative risk aversion is without loss of generality: in equilibrium only the marginal rate of substitution μ between consumption in a Low state and consumption in a High state needs to be estimated – the reverse marginal rate of substitution is $1/\mu$ and the marginal rates of substitution between consumption in two High states or two Low states are 1. Hence, only the marginal rate of substitution μ and the discount

factor β need to be estimated.

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

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

Each choice of instruments leads to a different test.

Our *first test* is based on a traditional instrument choice, going back to Hansen and Singleton (1983). We choose as instruments (i) the constant 1, (ii) lagged consumption growth, and (iii), lagged returns on the Tree T and (iv) lagged returns on the Bond B . Thus, we have 4 instruments, and hence have 8 moment conditions. Only two parameters β , γ need to be estimated, so we have 6 over-identifying restrictions. The idea behind GMM is to find values of the parameters that minimizes 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 over-identifying restrictions.²⁴ The necessary time series, of consumption growth and asset returns, were constructed by concatenating periods across all replications and all sessions, leaving out observations that would straddle two different replications, as we did for Table 9.

The top panel of Table 13 displays the results of the first test. We note three points:

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.

²⁴We implemented GMM using Matlab routines provided by Michael Cliff.

Table 13: 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 & 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)

However, it seems fair to say that the results of this test are misleading. In particular, GMM produces the estimate that the representative agent is risk-neutral or even slightly risk-loving – but the data we have *not* used in the test clearly show that subjects are risk averse: the Tree is cheaper than the Bond, and participants smooth consumption both across states and across time.²⁵ Closer inspection of the data suggests why GMM produces the peculiar estimate of risk aversion: average returns on the securities are positive (reflecting the significant discount rate) but the average return on the Tree (12.8%) is below that of the Bond (15.9%). GMM can only reconcile this perverse ranking of expected returns by assuming that the representative agent is risk-loving, and hence produces an incorrect estimate of risk aversion.

In looking more closely, we might view the χ^2 GMM test of over-identifying restrictions as suspect. Two instruments – the lagged returns on the Tree and Bond – are “weak”, in the sense that they are uncorrelated (even independent) over time, both with themselves and with consumption growth.²⁶ Hence these moment conditions do not provide additional restrictions beyond the ones imposed by the moment conditions constructed with the constant as instrument. Effectively, the number of degrees of

²⁵This finding is reminiscent of that in Asparouhova (2006), which studied a competitive market for loans under adverse selection. Standard structural estimation using the Rothschild-Stiglitz equilibrium pricing model failed to reject the model even when the model was obviously false, and yielded parameter estimates that were significantly different from the truth.

²⁶The details of the calculation can be obtained from the authors upon request.

freedom in the χ^2 test is not 6, but only 2.

To determine the impact of these weak instruments, we ran a *second test*, re-estimating the model with only the constant and lagged consumption as instruments. The second panel of Table 13 displays the results. The model fails to be rejected at an even higher p -level, the estimation of the discount factor β is nearly the same, but the estimation of risk attitude is even more risk loving.

However, GMM tests with traditional instruments do not exploit all restrictions of the model. In particular, expected returns are predicted to be different across states, but consumption growth can only capture the *change* in the state, and not the realization of the state. In historical data from the field, consumption growth is readily observable (though there is a debate whether the right consumption series is being used), but not the state itself. Here, we are in control of the state, and hence, can use it as an instrument. This is particularly appropriate because of the finding that the discount of the price of the Tree relative to the price of the Bond is counter-cyclical, in contrast to the theoretical prediction. (As we shall discuss below, the equity premium – which is the analog in terms of returns – is pro-cyclical, in violation of the theory.) GMM may be able to pick up this perverse result and thereby the reject the model.

Consequently, in our *third test*, we replaced the constant instrument with two dummy variables, one that tracked the high state, and the other one tracking the low state. We kept the remaining instrument, the consumption growth. In total, this gives three instruments and as such generated 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 13. 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 ($p = 0.001$).

Further inspection sheds light on why the GMM test (which is based on moment restrictions on asset returns) rejects. At the estimated parameters, most moment conditions fit tightly. However, the moment condition involving the Tree return and using the dummy variable for the High state does not fit well. And indeed, the average return on the Tree in the High state is lower (at 12.8%) than that of the Bond (at 15.9%). Thus, the equity premium in the High state is *negative*, and this can only be fit with a negative risk aversion coefficient. In contrast, in the Low state, the ranking of returns is consistent with risk aversion: 17.8% for the Tree and 16.1% for the Bond. GMM estimates a positive risk aversion coefficient, allows it to fit well the moment

conditions in the Low state for both assets, as well as the moment condition for the Bond in the High state (which is lower than in the Low state, consistent with the theoretical prediction). Notice also that the equity premium is pro-cyclical, contrary to the theory (but in line with the counter-cyclical nature of the discount of the Tree price against the Bond price).

7 Conclusion

Over the last thirty years, the Lucas model and the ideas that underlie it have become the central theoretical models through which scholars of macroeconomics and finance interpret the real world. Despite this, little is known about the true relevance of the Lucas model and confidence in the model has certainly 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 the model in the laboratory. Of course, it is a long way from the laboratory to the real world. There are many features of the real world that 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 experiment. As we have noted in the Introduction, 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 findings display 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, risky assets yield a substantial premium over riskless assets. Our experimental findings also display features that are at odds with other predictions of the Lucas model: prices display excess volatility, standard tests reject the stochastic Euler equations. Interestingly, the latter features are precisely those that much of the literature has attempted to explain in terms of real-world “frictions” or deviations from the basic model – frictions and deviations that are entirely absent in the laboratory environment.

The dimension in which our experimental findings diverge most sharply from the predictions of the Lucas model is in the ability of subjects to predict future prices. The Lucas model presumes that the representative agent perfectly forecasts future asset prices – but the subjects in our experiments do not do so. Given the length of the time

series that subjects can observe and the amount of noise, this is not surprising, Indeed, it would require something like 80 periods for subjects to learn how Tree prices correlate with fundamentals (while still being unable to discern an effect of fundamentals on Bond prices). 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, it must also be kept in mind that in many ways we have given our subjects a much simpler problem than they face in the real world. In our experiment, subjects *are told* the true dividend process – in the world they would have to learn it. In our experiment the true dividend process is *precisely stationary* – in the world it is not. (Indeed, it might be argued that the world is not stationary at all – that stationarity is just an assumption imposed on a model which would otherwise be completely intractable.) And finally, in our world there is a great deal more aggregate consumption risk than in the real world. 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 real world it is (using the Mehra and Prescott (1985) estimates of U.S. data) only about 1.08. Thus, in many dimensions, we give the Lucas model the best possible chance to succeed – and (in this dimension) it does not. Nevertheless 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. As we have noted, given the noise it would require 80 periods for subjects to learn about the true relation between prices and fundamentals; to obtain replications of 80 periods we would need to choose the continuation probability β uncomfortably close to 1. Perhaps more importantly, a replication of 80 periods would last at least 4-5 hours (following the very substantial initial time for training etc.); carrying out an experimental session of such a duration would be very difficult.

However, the fact that subjects' forecasts of future prices are wrong, and that this appears to have very significant effects, despite the fact that the errors do not appear to be far wrong – current forecasts are not far from future realizations – is important in itself. If perfect forecasting is difficult in our idealized laboratory environment, it is absolutely impossible in the real world; if small errors in forecasts can have big effects in the laboratory, perhaps they can also have big effects in the real world.

At various points we have made statements to the effect that subjects smooth consumption and insure dividend risk but they do not insure price risk. Subjects may not insure price risk, because either they don't perceive it, or they don't think they can insure it. We think it is the latter. The former would be credible if the theory were

right, because indeed, the closed-form equation for prices has only dividends in it, so there is no price risk (separate from dividend risk). If subjects did perceive price risk, we may be over-estimating risk aversion when fitting the Euler equations, because these equations assume that price risk derives from dividend risk, instead of being distinct from dividend risk. As such, we should refrain from drawing strong conclusions from the estimate of the coefficient of relative risk aversion, which is 0.16. (One should always be careful in drawing conclusions from estimates of parameters of models that are rejected.)

Our experimental findings also illustrate that it is dangerous to extrapolate from the individual to the market. As in our static experiments (Bossaerts, Plott, and Zame, 2007b), 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. (In particular, that some – perhaps many – individuals appear to display behavior that conforms to prospect theory (Barberis, Huang, and Santos, 2001) does not mean that markets would reflect that behavior in any obvious way.) The “laws” of the (financial) system are different from those of its parts.

The idea of looking at experimental findings using the methodology typically used to study historical field data was suggested to us by the pioneering work of Asparouhova (2006). Here, as there, we think it yields interesting insights: because much information is missing from historical field data, the analysis of such data may yield misleading conclusions. In this case (as in Asparouhova (2006)), unless one uses the “correct” instruments, the “usual” methodology, when applied to price data alone, fails to reject the null – which it should, given all the data available. In our case, the experimental environment itself suggests which should be the “correct” instruments; outside the laboratory, it would surely be much harder to know which instruments to use.

Appendix: Instructions (Type I Only)

Web Address: filagora.caltech.edu/fm/

User name:

Password:

INSTRUCTIONS

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.

3. Examples (for illustration only)

Tables 1 and 2 give two sample examples of outcomes in a session. It is assumed that the session ends after the 6th period. Table 1 shows the 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.

Table 2 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 $(\frac{5/6}{1-5/6})$ times the expected dividend (which is equal for both the tree and the bond), or $(5) * (0.5) = 2.50$. Trade is assumed to take place at 2.50. 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 \$15.00. It would have been \$13.00 if the session had terminated in period 5.

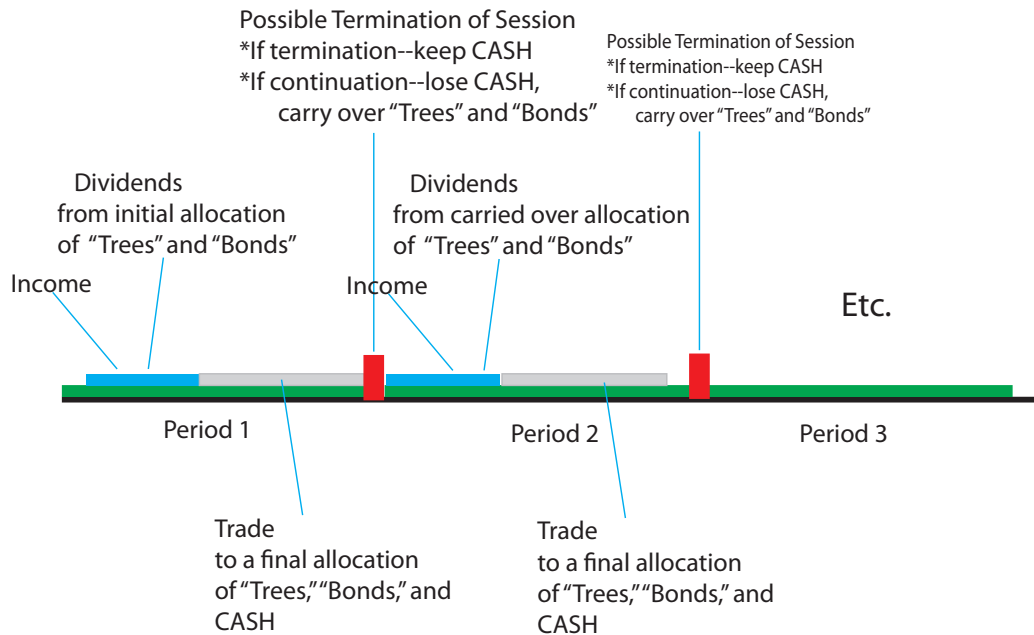
Table 1.

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	$\$0.5 \cdot 0 = 0$	$\$0.5 \cdot 0 = 0$	$\$0.5 \cdot 0 = 0$	$\$0.5 \cdot 0 = 0$	$\$0.5 \cdot 0 = 0$	$\$0.5 \cdot 0 = 0$
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

Appendix: Time Line Plot To Complement Instructions



References

- Klaus Adam, Albert Marcet, and Juan Pablo Nicolini. Stock market volatility and learning. Working paper, 2012.
- Elena Asparouhova. Competition in lending: Theory and experiments. *Review of Finance*, 10:189–219, 2006.
- 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/2696442>.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):pp. 637–654, 1973. ISSN 00223808. URL <http://www.jstor.org/stable/1831029>.
- Peter Bossaerts and William Zame. Risk aversion in laboratory asset markets. In J. Cox and G. Harrison, editors, *Risk Aversion in Experiments*, volume 12 of *Research in Experimental Economics*. JAI Press, 2008.

- Peter Bossaerts and William R. Zame. Asset trading volume in infinite-horizon economies with dynamically complete markets and heterogeneous agents: Comment. *Finance Research Letters*, 3(2):96 – 101, 2006. ISSN 1544-6123. doi: DOI: 10.1016/j.frl.2006.01.001. URL <http://www.sciencedirect.com/science/article/pii/S1544612306000031>.
- Peter Bossaerts, Charles Plott, and William Zame. Prices and portfolio choices in financial markets: Theory, econometrics, experiment. *Econometrica*, 75:993–1038, 2007a. Mimeo.
- Peter Bossaerts, Charles Plott, and William Zame. Prices and portfolio choices in financial markets: Theory, econometrics, experiments. *Econometrica*, 75:993–1038, 2007b.
- Colin F. Camerer and Keith Weigelt. *Research in Experimental Economics*, chapter An asset market test of a mechanism for inducing stochastic horizons in experiments, pages 213–238. JAI Press, 1996.
- John H. Cochrane. *Asset Pricing*. Princeton University Press, 2001.
- Sean Crockett and John Duffy. An Experimental Test of the Lucas Asset Pricing Model. 2013.
- Darrell Duffie and Chi-Fu Huang. Implementing arrow-debreu equilibria by continuous trading of few long-lived securities. *Econometrica*, 53(6):1337–1356, 1985.
- Kenneth B. Dunn and Kenneth J. Singleton. Modeling the term structure of interest rates under non-separable utility and durability of goods. *Journal of Financial Economics*, 17(1):27 – 55, 1986. ISSN 0304-405X. doi: DOI: 10.1016/0304-405X(86)90005-X. URL <http://www.sciencedirect.com/science/article/pii/0304405X8690005X>.
- Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 99(2):pp. 263–286, 1991. ISSN 00223808. URL <http://www.jstor.org/stable/2937681>.
- Eugene F. Fama. Efficient capital markets: Ii. *The Journal of Finance*, 46(5):pp. 1575–1617, 1991. ISSN 00221082. URL <http://www.jstor.org/stable/2328565>.

- John Geanakoplos and William R. Zame. Collateralized asset markets. Technical report, UCLA, 2007.
- Lars Hansen and K. J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91:249–265, 1983.
- Tarek A. Hassan and Thomas M. Mertens. The social cost of near-rational investment: Why we should worry about volatile stock markets. Working paper, 2010.
- Alejandro Hernandez and Manuel S. Santos. Competitive equilibria for infinite-horizon economies with incomplete markets. *Journal of Economic Theory*, 71(1):102 – 130, 1996. ISSN 0022-0531. doi: 10.1006/jeth.1996.0111. URL <http://www.sciencedirect.com/science/article/pii/S0022053196901118>.
- Charles Holt and Susan Laury. Risk aversion and incentive effects. *American Economic Review*, 92(5), 2002.
- Kenneth L. Judd, Felix Kubler, and Karl Schmedders. Asset trading volume with dynamically complete markets and heterogeneous agents. *The Journal of Finance*, 58(5):2203–2218, 2003. ISSN 1540-6261. doi: 10.1111/1540-6261.00602. URL <http://dx.doi.org/10.1111/1540-6261.00602>.
- Timothy J. Kehoe and David K. Levine. Debt-constrained asset markets. *The Review of Economic Studies*, 60(4):865–888, 1993. doi: 10.2307/2298103. URL <http://restud.oxfordjournals.org/content/60/4/865.abstract>.
- Stephen F. LeRoy and Richard D. Porter. The present-value relation: Tests based on implied variance bounds. *Econometrica*, 49(3):555–574, 1981. ISSN 00129682. URL <http://www.jstor.org/stable/1911512>.
- David K. Levine and William R. Zame. Debt constraints and equilibrium in infinite horizon economies with incomplete markets. *Journal of Mathematical Economics*, 26(1):103 – 131, 1996. ISSN 0304-4068. doi: 10.1016/0304-4068(95)00745-8. URL <http://www.sciencedirect.com/science/article/pii/0304406895007458>.
 ;ce:title;Special Issue on Equilibrium with Incomplete Markets and an Infinite State Space;ce:title;.
- Robert E. Jr. Lucas. Asset prices in an exchange economy. *Econometrica*, 46(6): 1429–1445, 1978. ISSN 00129682. URL <http://www.jstor.org/stable/1913837>.

- Hanno N. Lustig and Stijn G. Van Nieuwerburgh. Housing collateral, consumption insurance, and risk premia: An empirical perspective. *The Journal of Finance*, 60(3):pp. 1167–1219, 2005. ISSN 00221082. URL <http://www.jstor.org/stable/3694924>.
- Michael Magill and Martine Quinzii. Infinite horizon incomplete markets. *Econometrica*, 62(4):pp. 853–880, 1994. ISSN 00129682. URL <http://www.jstor.org/stable/2951735>.
- Burton G. Malkiel. *A Random Walk Down Wall Street*. W.W. Norton, New York, NY, 1999.
- Ian Martin. The lucas orchard. Working Paper 17563, National Bureau of Economic Research, November 2011. URL <http://www.nber.org/papers/w17563>.
- Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145 – 161, 1985. ISSN 0304-3932. doi: DOI: 10.1016/0304-3932(85)90061-3.
- Robert C. Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183, 1973a. ISSN 00058556. URL <http://www.jstor.org/stable/3003143>.
- Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):pp. 867–887, 1973b. ISSN 00129682. URL <http://www.jstor.org/stable/1913811>.
- Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68:1281–1292, 2000.
- Roy Radner. Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica*, 40(2):289–303, 1972. ISSN 00129682. URL <http://www.jstor.org/stable/1909407>.
- Richard Roll. A critique of the asset pricing theory’s tests part i: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2):129 – 176, 1977. ISSN 0304-405X. doi: DOI: 10.1016/0304-405X(77)90009-5. URL <http://www.sciencedirect.com/science/article/pii/0304405X77900095>.
- Bryan R Routledge and Stanley E. Zin. Generalized disappointment aversion and asset prices. *Journal of Finance*, In press, 2011.

Paul A. Samuelson. Proof that properly discounted present values of assets vibrate randomly. *The Bell Journal of Economics and Management Science*, 4(2):pp. 369–374, 1973. ISSN 00058556. URL <http://www.jstor.org/stable/3003046>.

Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review*, 71(3):421–436, 1981. ISSN 00028282. URL <http://www.jstor.org/stable/1802789>.

Vernon L. Smith, Gerry L. Suchanek, and Arlington W. Williams. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151, 1988. ISSN 00129682. URL <http://www.jstor.org/stable/1911361>.

Figure 1: Snapshot of the trading interface. Two bars graphically represent the book of the market in Trees (left) and in Bonds (right). Red tags indicate standing asks; blue 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 non-essential.

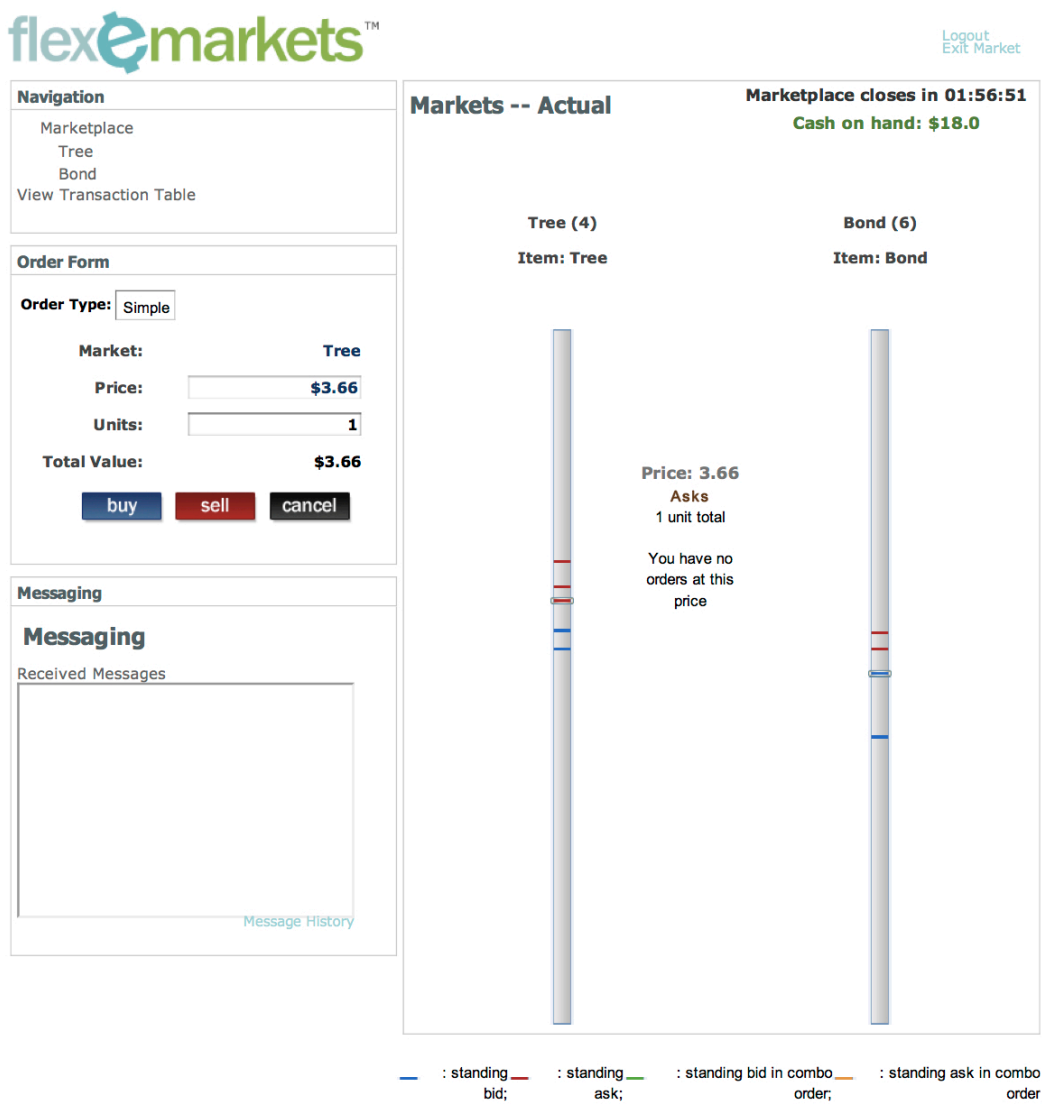


Figure 2: Time series of Tree (solid line) and Bond (dashed line) transaction prices; averages per period. Session numbers underneath line segments refer to Table 3.




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

