

‘Lucas’ In The Laboratory*

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

August 10, 2014

ABSTRACT

This paper studies the Lucas asset pricing model in a controlled setting. Participants could trade two long-lived securities in a continuous open-book trading system. Key experimental design features allowed us to emulate the stationary, infinite-horizon setting of the model and to incentivize participants to smooth earnings (consumption) across periods. Consistent with the predictions of the Lucas model, prices were aligned with consumption betas, they co-moved with aggregate dividends, and more strongly so when risk premia were higher. Trading significantly increased payoffs relative to autarky. Nevertheless, as in field data, prices were excessively volatile: aggregate dividends explained at most 18% of price changes, and residual price risk amounted to noise. The noise corrupted traditional GMM tests of the model. Choices displayed substantial heterogeneity, to the extent that no individual choice could be considered to be representative for pricing.

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

[†]University of Utah

[‡]University of Utah and University of Melbourne

[§]Singapore University of Technology and Design

[¶]UCLA

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

Here we study multi-period, multi-security asset pricing in a controlled setting and evaluate the outcomes against the predictions of the Lucas model. We find that consumption beta ranks security prices in cross section, as predicted by the model, but not security returns over time. We also find, as the model predicts, that fundamentals (aggregate consumption) drive changes in prices over time, implying (because of our parametrization) significant predictability, in violation of simple accounts of the Efficient Markets Hypothesis (EMH). Nevertheless, as in the field, the theory is rejected because the prices are excessively volatile: fundamentals explain at most 18% of price changes, while the theory predicts that fundamentals should explain 100% of price changes.

Since incentives are controlled and all choices are observed in the lab, we are able to study welfare. Across participant types (distinguished by initial holdings and period-income fluctuations), we discover substantial Pareto improvements from autarky, to the extent that consumption shares become statistically insignificantly different from constant, both over time and across states.² Thus, excess volatility does not appear to prevent improvements in welfare. However, we document considerable individual heterogeneity, so that no individual choice can be considered to be representative for pricing. Closer inspection of choices suggests that excess volatility is consistent with decision making based on subjects' (only mildly incorrect) prediction that prices do not change with fundamentals. The simulation of an economy populated with agents endowed with such beliefs leads to market outcomes that closely resemble the outcomes of our experiment.

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

From a methodological point of view, our novel experimental design addresses several difficult issues with laboratory testing of the Lucas model, such as the need to generate a stationary, infinite-horizon setting, and to make period-consumption perishable. Prices in our experiment do not “bubble,” a finding that is in contrast with the vast literature on

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

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

laboratory bubbles, and which we attribute to our design featuring incessant incentives to trade. This echoes the opinion of Crockett and Duffy (2013).

We proceed as follows. Because experimental tests of asset pricing theory are still rare, we first provide extensive motivation and then discuss the challenges one encounters when attempting to test the Lucas model in a controlled setting. We then present the Lucas model within a stylized version of the framework we created in the laboratory. Section IV provides concrete details of the experimental design. Results are discussed in Section V. Section VI uses the data from the laboratory to investigate the statistical properties of the GMM tests with which historical data from the field have traditionally been analyzed. Section VII discusses potential causes behind the excessive volatility of asset prices observed in our laboratory markets. Section VIII concludes. Appendix I discusses the theory in greater detail; Appendix II presents an illustrative numerical example.

I. Why Experimental Tests of The Lucas Model?

Controlled experimentation with markets is not standard methodology in empirical finance, so we first must address why we think our exercise has value. We think that experimentation itself should not be an issue. Controlled experiments are the foundation of science. That experimentation is still rare in finance must have to do with the difficulty of designing experiments that are informative. We would argue that whether experiments are “informative” is not only a question of *direct* relevance for ‘real world’ phenomena. The goal of experiments is, in the first place, to verify the veracity of theory in a setting where confounding factors are eliminated as much as possible. This is important because the model might fail to fit “real world” data because additional forces are at work, which the base version of the model does not account for. But if we cannot verify the veracity of the theory in a setting where confounding factors are eliminated – if the most basic version of the model indeed fails in some important dimension in a setting most conducive for it to work – then what is the point of using it or its extensions to explain the “real world”?

This is well understood in the sciences, and used even in disciplines that, like finance, cannot experiment at a “real world” scale, such as astronomy. To quote the late Hannes Alfvén, astrophysicist (and Nobel Prize laureate in Physics): “We have to learn again that science without contact with experiments is an enterprise which is likely to go completely astray into imaginary conjecture.” In our setting, Alfvén’s warning is that there is a real danger that the Lucas model is reduced to a convenient set of equations that one can fit to historical return data from the field, and that successful fit is considered proof that the theory works and unsuccessful fit is considered proof that the theory does not work. As an analogy, and at the risk of sounding pedantic, consider a (once well accepted!) theory of physics that claims that objects that are heavier fall faster, and that the link between speed and weight is to be determined empirically. With a suitable parameterization, this theory will fit “real world” data as well as most theories fit financial data. Yet we know

that the theory is plain false. How? Because of controlled experiments, where one of the confounding factors, friction, is suitably controlled.

We report that, even in its most basic setting, the Lucas model fails in one important respect, namely, the assumption of perfect foresight, and that this matters a lot for prices (it generates excess volatility), but less for allocations (substantial Pareto improvements do emerge). Since all extensions, with the exception of Adam, Marcet, and Nicolini (2012), rely on perfect foresight of the mapping from states to prices, the models can be used as a basis for understanding the “real world” only with this caveat in mind.

Of course, our experiment might be deemed to be meaningless. Indeed, the theory might be seen to fail because we ran the wrong experiment. (Galileo could have thrown, not two stones, but a hammer and a feather, and concluded that weight mattered for falling objects.) This is why we shall spend much attention detailing results that are at first not of much relevance to scholars who study historical asset returns from the field. These results (concerning volume relative to total asset supplies, relative pricing, movement of prices with fundamentals, relation between risk premia and co-movement of prices and fundamentals, Pareto improvements in final allocations, etc.) are in line with the Lucas model, and make us confident that the excess volatility finding that ultimately lead to rejecting the Lucas model is not just because of bad experimentation.

As such, our experiment should be seen in the first place as an attempt to validate the theory behind investigation of “real world” data rather than being directly relevant for the “real world”. We study to what extent the lens (model) through which “real world” data are being interpreted is valid. And we find it failing in one important respect.

With the exception of Crockett and Duffy (2013), the Lucas model has been tested exclusively on historical data from the field using statistical analysis³ or calibration⁴ of its core equilibrium restrictions, the stochastic Euler equations.⁵ It is fair to say that the general conclusion from these tests is that the model fails. One could wonder why bother testing in the laboratory something that is “obviously wrong” in the field? This is where we argue that experimental testing does have something important to add. It is precisely *because* the model fails on historical data that one wants to test whether it is true using controlled experiments. The fact that the stochastic Euler equations do not fit a given set of historical data from the field does not make the model invalid. Observation in the field is inevitably incomplete and it may very well be that one is making the wrong measurements (e.g., consumption of durable goods is left out or modeled incorrectly⁶). And as mentioned before, there may be forces at play that the theory abstracts from and are ignored in the empirical analysis (e.g., transaction costs). Or the theory just does not apply because an important assumption is violated (e.g., allocations are not Pareto efficient, and hence a representative agent does not exist). Etc.

³The seminal paper is Hansen and Singleton (1983).

⁴The seminal paper is Mehra and Prescott (1985)

⁵Analyzed in depth in, among others, Cochrane (2001).

⁶Dunn and Singleton (1986) test a version of the model where consumption goods provide services for two periods only.

Here, we propose an experimental design that represents a realistic setting yet with minimal complexity. Real people trade for real money in real markets, and their task encapsulates the two key goals of trading in the Lucas model, namely, diversification (across risky securities) and smoothing (of consumption over time). Our design accommodates agent heterogeneity (human subjects do exhibit differences in, e.g., attitudes towards risk), and imposes endowments that are nonstationary in order to induce trade beyond the first period. Of course, one cannot expect the Lucas model to work in all dimensions, but as previous experimental work on asset pricing theory has demonstrated (Bosschaerts, Plott, and Zame, 2007a), deviations from model predictions should be informative about facets of the theory that need to be improved and how improvements could be accomplished.

Our experiment does not provide a faithful replication of the field – nor should it. As is usual and proper in experiments, our goal is to eliminate the many confounding factors that stand in the way of evaluating the merits of a theory when using field data.

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

There is one important difference between the stationary-in-levels and stationary-in-growth models. Stationary-in-levels models readily generate equilibrium phenomena such as counter-cyclical equity premia (we shall do so here too) without having to appeal to counter-cyclical risk aversion (Routledge and Zin, 2011). Assuming standard preferences such as power utility, when the economy is stationary in levels, wealth is higher in a good state than in the bad state (states are defined as levels of wealth). As a result, the curvature of the utility function is lower in a high state than in a low state, so risk premia are lower in a high state than in a low state. One thereby obtains counter-cyclical equity premia. When the economy is stationary in growth, however, the state can be good (consumption growth was high leading into the period) or bad (consumption growth was low) with probabilities that are independent of wealth levels. (Notice that here states are defined in terms of growth, not wealth levels.) Since utility curvature only depends on wealth levels, and the state of the economy is unrelated to wealth levels, one cannot have counter-cyclical equity premia. To nevertheless obtain counter-cyclical equity premia, one could introduce a reference point that is low when the state is good and high when the state is bad (Campbell and Cochrane, 1999). This way, the curvature of the utility function not only depends on wealth, but also explicitly on the state. Counter-cyclical equity premia emerge naturally.

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

In the laboratory, inspection of an anomaly may help identify its origins. It may reveal the blind spots of the theory, and therefore help clarify why the theory fails to explain the field. Prices in the Lucas model are solely determined by (“are measurable in”) the fundamental risk in the economy, and price risk is equivalent to fundamental risk. In our experiments, there is only one source of fundamental risk, namely, aggregate consumption.⁷ Yet a large fraction of the price changes we observe is independent of changes in fundamentals. The existence of this apparent residual risk, distinct from fundamental risk, is inconsistent with the Lucas model. Within the model any such residual risk is eliminated by agents’ perfect knowledge of how prices change with fundamentals, i.e., by their perfect foresight. We not only observe this residual risk, which we shall refer to as *residual price forecasting risk*,⁸ but we document that it is large in magnitude: fundamentals explain at most 18% of the variability of securities prices, so residual price forecasting risk accounts for the remaining 82%.⁹

With hindsight, it perhaps should have been obvious that imperfect foresight is the Achilles heel of the Lucas model. But in fact it did not seem so obvious. Starting with Plott and Sunder (1988), evidence from financial markets experiments with heterogeneous information have given overwhelming support to the view that markets can manage to produce perfect foresight equilibria, to the extent that markets are being promoted as forecasting devices (Arrow, Forsythe, Gorham, Hahn, Hanson, Ledyard, Levmore, Litan, Milgrom, Nelson, et al., 2008). We are beginning to understand the aspects of human nature that facilitate this remarkable ability to acquire perfect foresight (Bruguier, Quartz, and Bossaerts, 2010). There, however, the setting is one where uninformed participants have to extract insider information from the order and trade flows. In contrast, in our case equally informed

⁷In the field, one can argue that past empirical investigations may have missed important risk components; see, e.g., Heaton and Lucas (2000)

⁸Note that given the structure of the Lucas economy, the residual risk can only come through the agents’ imperfect price forecasts.

⁹There have been recent attempts to incorporate in the Lucas model uncertainty about the true nature of the evolution of fundamentals (Maenhout, 2004; Epstein and Wang, 1994). It deserves emphasis that the resulting models continue to assume perfect foresight (of prices given future states). As such, these models still do not accommodate residual price forecasting risk.

participants have to build, from the same flows, correct expectations of how future clearing prices relate to fundamentals.

As we shall see, the design we chose to test the Lucas model requires participants to follow quite sophisticated hedging strategies. One may conjecture that our subjects (students at Caltech, UCLA and the University of Utah) lack the financial sophistication to conceive of, let alone execute, these strategies. Nevertheless, there are plenty of examples in the experimental economics, psychology, and computational neuroscience literatures that humans are capable of quite complicated behavior towards uncertainty. Payzan-Le Nestour and Bossaerts (2014), for instance, document that subjects resort to complex Bayesian learning strategies in a six-arm restless bandit problem. Subjects' ability to trade on insider information that they infer from order and trade flows likewise demonstrates surprising sophistication (Plott and Sunder, 1988). As to hedging, Bossaerts, Meloso, and Zame (2013) show how subjects resort to complex hedging strategies in a two-period asset market setting.

Controlled experiments can also shed light on the plausibility of important features of a theory in the presence of a more compelling alternative. Such is the case with the Lucas model, where prices are tightly linked to fundamentals, and hence, to the extent that fundamentals are predictable (e.g., they exhibit periodic cycles), prices must be predictable as well. In contrast, original accounts of the Efficient Markets Hypothesis (EMH) stated that prices must not be predictable: prices must be a martingale, except for drift as compensation for risk (Samuelson, 1973; Malkiel, 1999). The idea behind EMH is that investors would trade to exploit the predictability and in the process eliminate it. In historical data from the field, many cases have been discovered where securities prices over longer horizons can be predicted with, e.g., dividend yield (Goyal and Welch, 2003), where securities prices exhibit cyclicity (mean reversion; Lo and MacKinlay (1988)), or where price evolution can be predicted for several months after specific events such as earnings surprises (Bernard and Thomas, 1989). Interestingly, these violations of EMH tend to be explained, not as confirmation of the Lucas model, but in terms of behavioral finance, which is to say, in terms of cognitive biases in investor decision making (Bondt and Thaler, 1985; Daniel, Hirshleifer, and Subrahmanyam, 1998; De Long, Shleifer, Summers, and Waldmann, 1990). Lucas' proposition is that securities prices *may be predictable* even in properly functioning markets that do not display any behavioral anomalies. Our experiments demonstrate that this is not only a theoretical possibility, but also eminently relevant to real financial markets. In the experiments, prices move with the aggregate dividend (though in an extremely noisy way). Because aggregate dividends are predictable in the sense that the aggregate dividend is expected to fall in periods with high aggregate dividend, and *v.v.*, prices are too. Significantly, the aggregate dividend is the *only* variable that predicts price changes in our experiments.

In summary, the Lucas model is what it says: a model. It abstracts from many factors that play a role in the field, in order to get at essential insights. Models are idealizations and the laboratory is an idealized environment in which to truly test the model. That is, the laboratory is the right environment for a first test of any theory. We need such a test

because the model is being used as a lens with which historical data from the field are interpreted, and this lens is to be evaluated on its veracity, through controlled experiments.

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

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

Let us start with the last challenge, the existence of a representative agent. Unless agents are identical, which seems hardly more likely in the laboratory than in the field, the representative agent is only an equilibrium construct, and not a testable assumption/prediction. Fortunately for us, the heterogeneous agent version of the Lucas model yields predictions that are qualitatively no different than the predictions of the representative agent model (although they arise in a different way) and are testable in the laboratory environment. Pareto optimality plays a central role here. In the representative agent model, Pareto optimality is tautological – there is after all, only one agent. In the heterogeneous agent model, a representative agent can be constructed – but only if it is assumed that the result of trade is a Pareto optimal allocation – which is not guaranteed – and the particular representative agent that is constructed depends on the particular Pareto optimal allocation that obtains through trade.

For the market outcome to be Pareto optimal would seem to require that the market reach a Walrasian equilibrium, 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 *perfectly foresee* future prices (Radner, 1972), and that agents can employ investment strategies that exhibit the *hedging* features that are at the core of the modern theory of both derivatives pricing (Black and Scholes, 1973; Merton, 1973b) and dynamic asset pricing (Merton, 1973a).

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

(i) the time horizon is infinite and that agents discount the future, (ii) agents prefer to smooth consumption over time, and (iii) the economy is stationary. Meeting these challenges requires a novel experimental design. We deal with the infinite horizon as in Camerer and Weigelt (1996), by introducing a random ending time determined by a constant termination probability.¹⁰ More precisely, the experimental market unfolds period by period, and each such period can be terminal with the said probability. Termination uncertainty resolves at the conclusion of each period, after subjects have established their securities and cash holdings for that period. Upon termination, all securities expire worthless.

We provide an incentive for participants to smooth consumption by emulating perishability of consumption in each period: at the end of every 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 payout accounts (and hence "consumed").

Stationarity of the laboratory economy might seem evident, given the stationarity of the dividend process. However, stationarity of the termination probability does presents a severe difficulty. If an experimental session lasts for, say, 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 that the termination probability must be higher. However, if as a result subjects believe that the termination probability is not constant, a random ending time would induce a non-constant discount factor – and very likely induce different discount factors across subjects, because they may each imagine a different way of ensuring that the last period occurs before the session ends. To treat this challenge, we introduce a novel termination rule.¹¹ We elaborate on the rule in Section IV. Here we only mention that it relies on the assumption of stationarity of dividends in levels (and not in growth) and on the small (only two) number of possible dividend realizations.

In parallel work, Crockett and Duffy (2013) also study an infinite horizon asset market in the laboratory, but their experimental approach and purpose are different from ours. First, there is no risk in their setting. Our experiment includes risk, and hence, allows us to study the interplay between the two core drivers of the Lucas model, namely, risk avoidance (through diversification) and inter-temporal consumption smoothing. Second, Crockett and Duffy (2013) induce a preference for consumption smoothing by imposing a

¹⁰As 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).

¹¹One complication that is often raised as an issue for testing models like Lucas' in the laboratory is that they require participants to be risk averse. Consistent with evidence on individual decision making (Holt and Laury, 2002), prior experimental analysis of asset pricing models has demonstrated overwhelmingly that prices reflect risk aversion as theorists know it – decreasing marginal utility. See, e.g., Bossaerts, Plott, and Zame (2007a); Bossaerts and Zame (2008). This constitutes an anomaly only if one insists that participants evaluate all losses and gains in relation to present value of lifetime wealth. Rabin (2000) has cautioned us about the use of a single utility function to represent preferences over all ranges of wealth.

schedule of final payments to participants that is non-linear in period earnings. We take a different approach, and induce preference for consumption smoothing by paying only for one carefully chosen period (the last one), forfeiting payments in all periods that end not being terminal. Third, the focus in Crockett and Duffy (2013) is on the contrast between pricing in a treatment where there are incentives to trade because of demand for consumption smoothing and pricing in a treatment where there are no such incentives because there is no demand for consumption smoothing. Crockett and Duffy (2013) thus sheds light on a long line of experimental work on asset price bubbles, starting with Smith, Suchanek, and Williams (1988), where there are incentives to trade only in initial stages. In contrast, here we are interested in aspects of the Lucas model that have generated controversy in studies of historical asset prices. Of course, for good experimental control, and in line with Crockett and Duffy (2013), our setting is one where incentives to trade remain present throughout. A final difference concerns the stationarity assumption of the Lucas model. Like us, Crockett and Duffy (2013) used random termination to induce discounting. But we also needed to ensure stationarity, which we obtain through our novel termination protocol.

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

We formulate an instantiation of the Lucas asset pricing model that is simple enough to be implemented in the laboratory and yet complex enough to generate a rich set of predictions about prices and allocations, even under very weak assumptions. In particular, we allow for agents with different preferences and endowments (of assets and time-varying income streams), and we make no assumptions about functional forms, but still obtain strong and testable implications for individual consumption choices and trading patterns and for prices.

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

To create an environment suitable for the laboratory setting, we use a formulation that necessarily generates a great deal of trade; in our formulation, Pareto optimality (hence equilibrium) requires that trading takes place *every period*. This is important in the laboratory setting because subjects do not know the “correct” equilibrium prices (nor do we) and can only learn them through trade, which would seem problematic (to say the least) if trade were to take place infrequently or not at all. We therefore follow Bossaerts and Zame (2006) and treat a setting in which aggregate consumption is stationary (i.e. a time-invariant function of dividends) but individual endowments are not.¹²

¹²As Judd, Kubler, and Schmedders (2003) have 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.

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

A. The Structure Of The Laboratory Economy

We consider an infinite horizon economy with a single perishable consumption good in each time period. In the experiment, the consumption good is cash so we use the terms ‘consumption’ and ‘cash’ interchangeably here. How we make cash “perishable” and how we make the laboratory economy infinite-lived will be explained later. 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.¹³ We assume $d_T^H > d_T^L \geq 0$ and normalize so that the Bond and Tree have the same expected dividend: $d_B = \pi d_T^H + (1 - \pi) d_T^L$. Note that the dividend processes are stationary in levels. (In the experiment proper, we choose $\pi = 1/2$; $d_T^H = 1$, $d_T^L = 0$; $d_B = 0.50$, with all payoffs in dollars.)

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

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

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

where c_t is (stochastic) consumption at time t . We assume that the period utility functions u_i are smooth, strictly increasing, strictly concave and have infinite derivative at 0 (so that optimal consumption choices are interior), but make no assumptions as to functional forms.

¹³Lucas (1978) assumes that a *Tree* and a *one-period bond* are available; we use a *consol bond* simply for experimental convenience. The two formulations generate the same predictions with respect to consumption and equivalent predictions with respect to prices.

Note that agent endowments and utility functions may be heterogeneous but that all agents use the same constant discount factor β , which we induce to equal $5/6$.

In each period t agents receive dividends from the Bonds and Trees they hold, trade their holdings of Bonds and Trees at current prices, use the proceeds together with their endowments to buy a new portfolio of Bonds and Trees, and consume the remaining cash. How exactly agents buy and sell in our laboratory economy will be explained later on. Here, we follow the theory and assume that 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 (“perfect foresight”). This is not quite enough for equilibrium to be well-defined because it does not rule out the possibility that agents acquire more and more debt, delaying repayment further and further into the future – and never in fact repaying it. In order 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.¹⁴ (In the experimental setting, we forbid short sales so debt is necessarily bounded.)

B. Predictions

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

Despite allowing for heterogeneity and without making any assumptions about functional forms for preferences, the theory makes testable *quantitative* predictions about individual consumptions (which, in the experiments, will be end-of-period cash holdings), prices and

¹⁴Lucas (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.

trading patterns. Some of these predictions take a particularly simple form when the specific parameters are as in our experiments. The pricing predictions should be entirely familiar in the context of the usual Lucas model with a representative agent having CRRA (constant relative risk aversion) utility. We go far beyond pricing because we formulate predictions about allocations as well.

Prediction 1. There is trade in each period.

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

Prediction 2. The tree is always cheaper than the bond.

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

In an economy with heterogeneous agents, the preferences of the representative agent are hard to derive. These preferences may not even look like those of any individual agent. However, since each period we only have two possible states (the state σ is either high H or low L), things simplify dramatically.

To see this, fix an individual agent i ; write $\{c_i\}$ for i 's equilibrium consumption stream. Write i 's first-order condition for optimality:

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

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

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

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

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

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

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

Solving yields:

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

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

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

The important thing to note here is that:

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

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

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

Prediction 3. Asset prices are perfectly correlated with fundamentals.

From (3) it follows that:

$$p_A^H/p_A^L = \mu. \quad (4)$$

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

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

Prediction 4. The more the Tree trades at a discount relative to the Bond, the greater the difference of prices of both securities across states.

More precisely, cross-sectional and inter-temporal features of asset prices reinforce each other.¹⁵ The discount of the Tree relative to the Bond increases because risk aversion rises. As a result, the Tree and Bond prices move more extremely with fundamentals. Mathematically,

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

for $\sigma = H, L$ and $A = B, T$, with covariance computed based on sampling across cohorts of agents (economies), keeping everything else constant. "Everything else" means: initial endowments, private income flows, asset structure, outcome probabilities, as well as impatience β . Economies are therefore distinguishable at the price level only in terms of the risk aversion (embedded in μ) of the representative agent.¹⁶

Prediction 5. Expected returns vary across states.

From prices and dividends of the Tree and Bond as well as the state transition probabilities (from a Low state one moves to a High state or remains in the Low state with equal probability), one can readily compute the expected returns on the Tree and the Bond. Simple algebraic manipulation then allows one to express the difference across High and

¹⁵This insight forms the basis of recent GMM tests on field data, whereby parameters obtained from the cross-section of asset returns are used to verify intertemporal equilibrium restrictions. See Nagel and Singleton (2011).

¹⁶To 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.

Low states of the expected return on the Tree ($E[R_T|H] - E[R_T|L]$) and Bond ($E[R_B|H] - E[R_B|L]$) as follows:

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

Because $\mu > 1$, all terms in both expressions are negative, and hence, the expected return on both assets is higher in “busts” (when the state is Low) than in “booms” (when the state is High). Lucas (1978) emphasized return predictability as one of the natural equilibrium features of the model and contrasted this with the martingale version of EMH.¹⁷

Prediction 6. The equity premium is counter-cyclical.

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).¹⁸ Because asset prices and dividends are stationary, equity premia are stationary as well and depend only the state. Simple computations then allow us to explicitly derive the equity premia in the High and Low states. These computations rely on two facts: (i) the expected dividends are the same for both assets and equal to d_B , and (ii) for each asset A , $p_A^H = \mu p_A^L$. We obtain:

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

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

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

This difference is strictly negative (because $\mu > 1$), so the equity premium is counter-cyclical: it is lower in the High state than in the Low state. Note that counter-cyclicity

¹⁷Of course prices *do* form a martingale under the risk-neutral probabilities – the probabilities adjusted by marginal rates of substitution – but the risk-neutral probabilities are equilibrium constructions because marginal rates of substitution depend on equilibrium allocations.

¹⁸Mehra and Prescott (1985) use a slightly different model, with long-lived Tree and a one-period bond, and define the equity premium as the difference between the expected return on the risky security and that of the one-period bond.

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

Prediction 7. Agents smooth consumption over time.

This is an immediate consequence of Pareto optimality, the condition that guarantees the existence of a representative agent. Despite the fact that individual incomes fluctuate over time, in equilibrium agents smooth out this fluctuation by trading in securities. Income fluctuations do not impact available aggregate consumption, and hence, agents should be able to trade them away.

Prediction 8. Agents hedge fundamental risk.

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

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

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

In summary, our implementation of the Lucas model predicts that securities prices differ cross-sectionally depending on consumption betas (the Tree has the higher beta), while intertemporally, securities prices move with fundamentals (dividends of the Tree). The two predictions reinforce each other: the bigger the difference in prices across securities, the larger the intertemporal movements. These predictions would also follow from the original Lucas model. But in our setting, there is agent heterogeneity, and we can say more. Specifically, investment choices should be such that consumption (cash holdings at the end of a period) across states becomes perfectly rank-correlated between agents (or even perfectly correlated, if agents have the same preferences). Likewise, consumption should be smoothed across periods with and without income. Investment choices are sophisticated: they require, among others, that agents hedge price risk, by buying Trees when experiencing income shortfall (and selling Bonds to cover the shortfall), and selling Trees in periods of high income (while buying back Bonds).

For illustration, the Appendix provides explicit solutions for equilibrium prices, holdings and consumptions when taking the parameters as in the experiment and assuming that all agents display identical constant relative risk aversion, which we vary from 0.2 to 1.0. There is nothing special about these particular choices of risk aversion; we offer them solely for comparison purposes. We note that risk aversion in the range .2 – .5 is consistent with the experimental findings of Holt and Laury (2002) and Bossaerts and Zame (2008).

IV. Implementation In The Laboratory

As we have already noted, implementing the Lucas economy in the laboratory 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.¹⁹

So, while we will not literally have consumption every period, agents in our setting face the same optimization problem as if they actually had to consume every period and wanted to smooth consumption across periods. Thus, our experiment is designed so that, for subjects to optimize, they need to solve the same problem in the laboratory as in the theory. When choices for the two optimization problems are the same, the Lucas model predicts that prices will also be the same, because in the Lucas model, prices *derive* from choices, by means of the stochastic Euler equations.

It is less obvious how to resolve problem (c). The problem is not with the dividends and personal income but with the termination probability. In principle, simply announcing a constant termination probability should do the trick – because each period is equally likely to be terminal. However, if the probability of termination is in fact constant (and independent of the current duration) then the experiment could continue for an arbitrarily long time. In particular there would be a 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 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. Note that there should be no trade in the terminal period because assets perish at the end and hence are worthless. We nevertheless opened markets, to verify whether subjects understood this. Invariably, they did. In the

¹⁹Starting 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.

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

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

The series telescopes, so can be written recursively as in (1), which is the equation that applies to the infinite-horizon case. Hence, our two-period termination protocol produces pricing *as if* the economy were to continue forever.²⁰

In short, to utilize the two hour session time optimally, we ran as many replications as possible. Whenever a replication terminated and there was still time left in the session, we initiated a new replication. Thus, the termination protocol was only applied to at most one of the replications in a session. If a replication ended naturally (i.e., through the roll of the die) close to the 10-minute mark before end of a session, we did not start a new replication.

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 happens to be the terminal period; securities are desirable because they generate experimental earnings in future periods if the current period happens not to be the terminal period. The subjects had no difficulty grasping the essence of this task, and the simplicity allowed us to make instructions short. See the Appendix for sample instructions.

²⁰Our early-termination protocol relies on the specific setting we are in. Specifically, pricing in our two-period early-termination protocol is identical to that in an perpetual continuation of the random termination protocol because we are assuming that state drawing is independent over time. If we were to assume dependence, we could nevertheless alter the early termination protocol. For instance, if the state transition is markovian (meaning that the current state alone determines the transition probabilities) then a three-period early-termination protocol could be conceived, whereby the session is ended either in the first, second or third stage. The key is to be able to re-write pricing in this three-period world as an infinite series, and to show that this infinite series is identical to the pricing formula one obtains under perpetual continuation of the random termination protocol.

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 – borrowing – entails a promise to repay in the future. However, in the model, nothing enforces these promises; that they are kept is simply part of the definition of equilibrium. If nothing 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 penalties be? The rules governing experimentation with human subjects prevent us from forcing subjects to pay from their own pockets, and excluding subjects from further participation in the experiment would raise a host of problems following such 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 our theoretical analysis presumed 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 zero he could not sell it and if his consumption and portfolio were both zero he could not buy it. However, so long as agents do not bump up against the zero bound, our analysis remains correct. In the actual experimental data, the number of agents who bumped up against the zero bound was quite small.

The need to bar short sales explains why we use an instantiation of the Lucas economy in which the Bond is in positive net supply: risk tolerant subjects could merely reduce their holdings of Bonds rather than having to sell short (which 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.²¹ 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.²² Subjects quickly familiarized themselves with key aspects of trading in the open-book mechanism (bids, asks, cancellations, transaction determination protocol, etc.) through one mock replication of our economy during the instructional phase of the experiment. A snapshot of the trading screen is 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.

As explained above, 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 (2) of the replications, randomly chosen after conclusion of the experiment.²³ However, we do not view wealth effects as important in this context in any event, since there is no assumption that subject preferences are constant across replications within a given session – and they are certainly not constant across sessions, since the populations of subjects in different sessions were disjoint.

Finally, we allocated assets and income processes as follows.

- 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.

V. Results

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

²¹As we have noted earlier, we agree with Crockett and Duffy (2013) that frequent trading deters the formation of pricing anomalies such as bubbles.

²²Flex-E-Markets is documented at <http://www.flexemarkets.com>; the software is freely available to academics.

²³If a session ended with fewer replications we paid for multiples of some or all of the replications.

Table I
Summary data, all experimental sessions. Stars indicate sessions when last replication ended with a switch to the two-period termination protocol.

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

We now discuss to what extent each prediction is borne out in the data. We first re-state the prediction (we keep the the same order as before) and then evaluate the evidence.

Prediction 1. There is trade in each period.

Table II lists average trading volume per period (excluding terminal periods during which there was no trade). Trading volume in Periods 1 and 2 is significantly higher, reflecting the need for agents to move from initial holdings to their steady-state holdings. In the theory, subsequent trade takes place only to smooth consumption across odd and even periods.²⁴ 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 3-4 minutes long, this means that roughly 10% of available securities were traded each period and that one transaction occurred roughly every 5 seconds.

Prediction 2. The tree is always cheaper than the bond.

Table III displays average period transaction prices; Table IV stratifies the results by the period's state (High if the dividend of the Tree was \$1; Low if it was \$0). Consistent with the theory, the Bond is priced above the Tree; the price differential averaged about \$0.50.

²⁴In 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 II
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

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

	Tree	Bond	Discount
	Price	Price	
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

We momentarily refrain from making statements about the statistical significance. We will do so later when we discuss a proper modeling of the dynamics of (average-period) prices. Indeed, the time series exhibit peculiar properties, as we discuss next.

Prediction 3. Asset prices are perfectly correlated with fundamentals.

Prices moved with fundamentals, but correlation is *far from* perfect. Figure 2 shows a plot of the evolution of (average) prices over time, arranged chronologically by experimental

Table IV
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

sessions (numbered as in Table I); replications within a session are concatenated.²⁵ The plot reveals that prices are volatile. In theory, prices should move only because of variability in economic fundamentals, which in this case amounts to changes in the dividend of the Tree. Prices should be high in High states, and low in Low states, and these levels should be determined only by the state, not by, e.g., passage of time, or whether the period is Odd or Even. In reality, a large fraction of price movements is unrelated to fundamentals; following LeRoy and Porter (1981) and Shiller (1981), we 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. That is, after accounting for changes in states, price changes are white noise.

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

Prediction 4. The more the Tree trades at a discount relative to the Bond, the greater the difference of prices of both securities across states.

The prediction is that the differential in prices between High and Low states should correlate positively with the difference between the Bond price and the Tree price, i.e., the discount of the Tree price relative to the Bond price. Correlation is to be taken across economies, i.e., across replications. Table V displays correlations of the average discount on the Tree price relative to the Bond price (regardless of state) and the average difference between prices of the Tree or of the Bond across states. Each observation corresponds to one replication, so

²⁵This means that the average prices of the first period of the next replication are plotted as if it were the continuation of the previous replication. The second session, for instance, has 13 observations; the first 4 observations pertain to the first replication, and the subsequent nine observations pertain to the second replication.

Table V
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)

there are 15 observations in total. Consistent with the theoretical prediction, the correlations are positive, though the estimate is insignificant for the Bond.

Formal Statistics for Predictions 2-3-4

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

However, closer inspection of the properties of the error term revealed substantial dependence over time, as evident from the Durbin-Watson test.²⁶ Further model specification

²⁶The Table below displays the original estimation results. Notice that the Durbin-Watson (DW) test generated values that correspond to $p < 0.001$, suggesting that the error term was highly autocorrelated, and hence, that the model was mis-specified.

Footnote Table

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

Explanatory 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 VI
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

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 VI.²⁷ 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), entirely driven by the actual sampling of the state. The autocorrelations of the error terms are now acceptable (comfortably within two standard errors from zero).

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

The excess volatility of prices is plainly evident from Table VI. Fundamentals (changes in the state) explain only 18% of the variability of the Tree prices ($R^2 = 0.18$); 82% of price variance is unexplained.²⁸ The situation is even worse for the Bond: 96% of the variance of Bond price changes is unexplained by changes in the state. The Lucas model instead predicts that fundamentals should explain *all* of the variability of asset prices. It 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 regressions

Notice though that the coefficient to the termination dummy is insignificant, suggesting that our termination protocol was indeed neutral, as predicted by the theory. This constitutes comforting evidence that our experimental design was correct.

²⁷We deleted observations that straddled two replications. Hence, the results in Table VI are solely based on intra-replication price behavior. Unfortunately, this deprived us from testing whether our session termination protocol was neutral. The original specification did allow us to perform this test, and as mentioned in the previous footnote, the results confirmed the neutrality.

²⁸We relate price changes to state changes using a linear model; however, because there are only two states, linearity is without loss of generality.

Table VII
Average returns across securities and states (High or Low aggregate dividend).

State	Tree	Bond	Equity Premium
High	12.8 (%)	15.9	-3.1
Low	17.8	16.1	1.7
Difference	-5.0	-0.2	-4.8
Average	16.1	16.0	0.1

in first differences show that, consistent with the Lucas model, fundamental economic forces are behind price changes – significantly so for the Tree. But at the same time, prices are excessively volatile, with no discernible drift.

Prediction 5. Expected returns are time-varying.

The theory states that the expected returns on the two securities should vary depending on the state. Specifically, expected returns are higher in the Low than in the High state, which is to say that expected returns are counter-cyclical. Table VII provides the evidence. The results are consistent with the theory. The counter-cyclicity of expected returns implies that returns *are* predictable, in violation of traditional accounts of EMH.

Prediction 6. The equity premium is counter-cyclical.

Table VII demonstrates that the equity premium is counter-cyclical. But the theory also requires the equity premium to be positive, while it is not in “boom” periods: the average return on the Tree is below that of the Bond when the state is High. This return anomaly is in sharp contrast with the general correctness of price levels. This paradoxical situation may be understood in view of the noise in prices, and since returns are ratios of prices (and dividends), the noise is amplified in the returns. As a result, the fit of the Lucas model is much worse when one considers returns rather than price levels.

Prediction 7. Agents smooth consumption over time.

In the Lucas equilibrium, consumption choices are Pareto optimal. This means, in particular, that agents of both types should trade to holdings that generate high consumption in High states, and low consumption in Low states. Table VIII displays the average amount of cash (consumption) per type in High vs. Low states.²⁹ Consistent with the theoretical prediction, consumption is positively correlated across Types.

To gauge the significance of this finding, Table VIII also displays, in parentheses, the consumption (cash) levels that agents could have reached if they were not allowed to trade.

²⁹To compute these averages, we ignored Periods 1 and 2, to allow subjects time to trade from their initial holdings to steady state positions.

Table VIII

Average consumption (end-of-period cash, in dollars) across states (High or Low Tree dividend) and across periods (Odd/Even), stratified by participant Type (autarky 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. For choices to be Pareto efficient, consumption shares should be independent of State and Period provided the representative agents for the two Types have homothetic utility.

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

These are the consumption levels under autarky. Note that autarky consumption levels are *anti-correlated*. Through trading, the average Type I and Type II agents managed to move correlation of their consumptions from negative to positive. This suggests economically significant Pareto improvements, consistent with the Lucas model.

Pareto optimality also implies that subjects should be able to perfectly offset income differences across odd and even periods. Table VIII demonstrates that our subjects 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), again implying substantial Pareto improvements.³⁰

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

³⁰Autarky consumption of Type II subjects is independent of the state, because Type II subjects are endowed with Bonds whose dividends are riskless; autarky consumption of Type I subjects depends on states because they are endowed with Trees, whose dividends are risky. We used the sequence of realized states across all the sessions to compute their autarky consumption.

accommodate differences in consumption shares across replications due to differences in drawing of the state in the first period and in subject cohort.

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

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

Prediction 8. Agents hedge fundamental risk.

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

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

This finding should put our trading volume numbers in perspective. As noted earlier, trade volume was substantial (see Table II). However, the volume of trade is much lower than predicted by the Lucas model. In the Appendix, we show that even for (constant) relative risk aversion as low as 0.2, average turnover should be about 20% of average holdings for the Tree, and over 50% for the bond. Keep in mind however that these predictions presume that subjects follow the sophisticated hedging strategy of Prediction 8 namely, to buy Trees when in need of cash. Without this strategy, the volume needed to smooth consumption is substantially lower. Fewer assets have to be traded since Tree purchases (required for hedging) need not be offset by extra Bond sales (required to smooth consumption).

Subject-Level Differences

There are substantial individual differences in portfolio choices. Table X 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 in periods with income shortfall (odd periods), effectively implementing the hedging strategy of the theory. Subject 5 is almost a mirror image of subject 7, though s/he did not resort to hedging. Subject 3 diversified across Trees and Bonds but likewise did not hedge because Tree holdings decreased in odd periods.

Table IX
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

The subject-level differences reported in Table X are not exceptional. The contrast between individual choices and choices aggregated over individuals of the same Type is stark. In terms of choice predictions, the theory “works” at the Type level, but not at the individual level and so one has to be careful extrapolating to phenomena at the market level (e.g., prices) from observing individuals singly. If we had considered Subjects 3 or 5 to be “representative,” and had predicted cross-sectional and temporal behavior of prices on the basis of their choices, the fit would have been poor. The situation is reminiscent of the cross-sectional variation in choices in static asset pricing experiments. There too, prices at the market level can be “right” (satisfy, e.g., CAPM) even if individual choices are at odds with the theory; see Bossaerts, Plott, and Zame (2007a).

Table X shows that subjects in certain periods held zero units of a security. If subjects followed equilibrium strategies and their risk aversion was not too high, this should not have happened; see Table XIII in the Appendix. The following provides some idea of how often subjects traded to the boundary. We exclude period 1, because as mentioned before, it often took more than one period for subjects to trade away from initial holdings (all subjects started with zero holdings of one of the two securities). Of the 1272 period-subjects (one subject per period), 19% held zero of the Tree, and 24% held zero of the Bond. Are these numbers high? This is difficult to say. Again, holdings of individual subjects are enormously variable and this variability reveals nothing about the success of the Lucas model.

Summary

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

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

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

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

VI. Using Laboratory Data To Assess The Statistical Properties Of GMM Tests Of The Lucas Model

As noted before, we (control and) observe much more in our laboratory environment than is possible in the field. However, an interesting and potentially revealing exercise is to ignore this additional richness, treat our laboratory data as if they were historical field data, and carry out the same kind of econometric tests that have been used for historical field data. Hence, let us consider only the times series of asset returns and aggregate consumption in our laboratory data. We ignore all the additional information (true dividend process, true realized state, individual choices, etc). Then let us use Generalized Method of Moments (GMM) to test whether the stochastic Euler equations are satisfied. 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) and test whether the Euler equations hold, while imposing that the representative agent “consumes” the aggregate cash each period.

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

It is our opinion that the theory explains many facets of the experimental results. However, as mentioned before, the behavior of asset returns is anomalous in one important dimension. Specifically, Table VII shows that 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. The corresponding Euler equation can only be fit with a negative risk aversion coefficient (assuming that β is set at its correct value, namely, 5/6). 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 will record the tension between the equity premium in the Low state and in the High state – as long as one uses the (true) state as an instrument so that the conditional average returns across the two states are contrasted. The GMM statistic should then reject the model.

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

The Euler equations are:

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

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

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

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, 8 moment conditions. Only two parameters, β and 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.³¹ 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 VI.

The top panel of Table XI 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.

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

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.³² 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 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 row of Table XI displays the results. The model still fails to be rejected (with an even higher p -level). And the estimation of risk attitude continues to be anomalous, because the point estimate remains in the risk-loving range, though insignificantly different from risk neutrality.

³¹We implemented GMM using Matlab routines provided by Michael Cliff.

³²The details of the calculation can be obtained from the authors upon request.

Table XI
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)

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

Consequently, in our third test, we 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 XI. We observe the following.

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

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

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

too low. The mis-estimation was solely due to lack of power when implementing the GMM test with traditional instruments.

VII. The Expected and the Anomalous

Viewed against the predictions of the Lucas model, our experiments generate findings that are expected – prices and individual consumption are correlated with fundamentals (aggregate consumption) and prices are ordered across securities and across states as predicted by the theory – and findings that seem anomalous – prices are excessively volatile and subjects do not hedge. Because volatile prices would seem to signal clearly the need to hedge, the co-existence of excess volatility and lack of hedging seems surprising. However, the particular kind of excess volatility that we see in the experimental data might well lead subjects to conclude that there is no opportunity to hedge.

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

Optimization against exactly correct beliefs leads exactly to the Radner equilibrium predicted by the Lucas model. It would seem that optimization against approximately correct beliefs should lead to something that approximates the Radner equilibrium predicted by the Lucas model. However, this is not necessarily so: because the price process is *endogenous*, beliefs about the price process can be approximately correct (i.e. close to the realized price process) and still very far from the price process predicted by the Lucas model. The same point has been made by Adam, Marcet, and Nicolini (2012), who used it to explain excess volatility in historical data.

On the basis of our experimental data, it seems quite plausible that agents expected prices to follow a martingale – as would be predicted by traditional accounts of EMH! – and *not* to co-move with economic fundamentals – as would be predicted by the Lucas model. 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, e.g., Bond prices follow a martingale would be hard to falsify even after eighty observations (see Table VI), which is an order of magnitude more observations than were available to subjects. The martingale belief is therefore a credible working hypothesis.

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

that agents then send demand schedules to the market; and that markets generate prices so that demand and supply are equal in that period. Of course, beliefs will be revealed to be wrong next period, so we are considering in this thought experiment only a kind of temporary equilibrium (Grandmont, 1977), 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 co-move with dividends, but very noisily – hence they would be excessively volatile.

To illustrate this thought experiment, we carried out simulations. There are two types of agents, endowed as in the experiment, and each type is represented by an agent with logarithmic utility. Agent beliefs (that prices revert to the levels of the previous period) are affected every period by an additive gaussian disturbance with mean zero and standard deviation \$0.40. Agents start out believing that the Tree will be priced at \$2.5 and the Bond at \$3. Figure 3 displays the evolution of prices and states in a typical simulation. Price dynamics are in line with those in the experiment – prices relate to states only in a noisy way. At the same time, agents turn out not to hedge residual price forecasting risk; they accommodate income shortfalls solely by selling Bonds and Trees. Still, their choices move substantially towards Pareto optimality: the consumption share of the Type I agent fluctuates 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 processes. This comes as a surprise because price predictions in the Lucas model are robust to small mistakes in beliefs about the dividend processes (Hassan and Mertens, 2010). Consequently, it seems that mistakes in beliefs can manifest themselves quite differently: when they concern the dividend processes, mistakes are dampened out; when they concern price processes, mistakes can create positive feedback and hence are amplified.

VIII. Conclusion

Over the last thirty years, the Lucas model of dynamic asset pricing has become the lens through which scholars of finance and macroeconomics interpret historical data coming from the field. Yet little is known about the true relevance of the Lucas model and confidence in the model has 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 a Lucas economy 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 experimentation. Models are idealizations and the laboratory is an idealized environment; if the models do not work in the laboratory, why should we expect them to work in the real world?

In our view, our experiments provide substantial support for the Lucas asset pricing model. Our experimental results display features that are consistent with the most important predictions of the Lucas model: prices move with fundamentals, agents trade assets to smooth consumption and insure against risk, more risky assets yield a substantial premium over less risky assets, and the equity premium is counter-cyclical. Our experimental results also display features that are at odds with the predictions of the Lucas model: prices display excess volatility, the equity premium is actually negative in “booms,” which lead standard tests to reject the stochastic Euler equations provided the right instruments are employed. Interestingly, excess volatility and rejections of the Euler equations are precisely those features of field data that much of the literature has attempted to explain in terms of “frictions” or deviations from the basic model. These frictions and deviations are entirely absent in our laboratory environment, suggesting that they may not be needed to explain failures of the model.

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 cannot and do not do so. Given the length of the time series that subjects can observe and the amount of noise, this is not surprising. It might be argued that, because our experiments are short (the longest replication is only ten 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 field they would have to learn it. In our experiment the true dividend process *is stationary*; in the field it is not. (Indeed, it might be argued that the world is not stationary at all, but that stationarity is just an assumption imposed on a model which would otherwise be completely intractable.) Thus, in many dimensions, we gave the Lucas model the best possible chance to succeed.³³

It would be of interest to know whether our subjects would eventually learn to make correct forecasts if they had many more observations. Unfortunately a design that provides enough observations seems quite impractical. We have noted that it would require eighty periods for subjects to learn about the true relation between prices and fundamentals; to obtain replications of eighty periods we would need to choose the continuation probability β uncomfortably close to 1. The high continuation probability may entice subjects not to take serious the trade-off between cash on hand and securities holdings. Perhaps more importantly, a replication of eighty 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 difficult.

³³We should add that in our experiment there is a great deal more aggregate consumption risk than in the field. In our experiment the ratio of aggregate consumption in the High dividend state to aggregate consumption in the Low dividend state is 1.50, while in the field it is (using the Mehra and Prescott (1985) estimates of U.S. data) only about 1.08.

However, the fact that subjects' forecasts of future prices are wrong, but not too wrong, 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 could also have big effects in the real world.

Mistakes in price forecasts cause prices to move beyond changes in fundamentals. If agents are aware that, besides fundamental risk, there is such residual price forecasting risk, then the equilibrium (Euler) equations are misleading. This is because the equations assume that there is no residual price forecasting risk. One would expect the agents themselves to be at least as averse to fundamental risk (about which they know a lot) as to residual price forecasting risk (about which they do not know much). Their investments, and hence, consumption choices, are likely to be affected at least as much by residual price forecasting risk. But the equilibrium equations explain consumption choices only in terms of fundamental risk, and consequently, there is a real danger that agents will look far more risk averse through those equations than they actually are.

Our experimental findings also illustrate that one should be cautious when extrapolating from the individual to the market. As in our static experiments (Bossaerts, Plott, and Zame, 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. The "laws" of the (financial) system may be different from those of its parts.

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

References

- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini, 2012, Stock market volatility and learning, Working paper.
- Arrow, Kenneth J, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D Nelson, et al., 2008, The promise of prediction markets, *Science* 320, 877.
- Asparouhova, Elena, 2006, Competition in lending: Theory and experiments, *Review of Finance* 10, 189–219.

- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.
- Bernard, Victor L., and Jacob K. Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium?, *Journal of Accounting Research* 27, pp. 1–36.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, pp. 637–654.
- Bondt, Werner F. M. De, and Richard Thaler, 1985, Does the stock market overreact?, *The Journal of Finance* 40, pp. 793–805.
- Bossaerts, Peter, Debrah Meloso, and William Zame, 2013, Dynamically-complete experimental asset markets, Available from author’s webpage.
- Bossaerts, Peter, Charles Plott, and William Zame, 2007a, Prices and portfolio choices in financial markets: Theory, econometrics, experiment, *Econometrica* 75, 993–1038, Mimeo.
- Bossaerts, Peter, Charles Plott, and William Zame, 2007b, Prices and portfolio choices in financial markets: Theory, econometrics, experiments, *Econometrica* 75, 993–1038.
- Bossaerts, Peter, and William Zame, 2008, Risk aversion in laboratory asset markets, in J. Cox, and G. Harrison, eds., *Risk Aversion in Experiments*, volume 12 of *Research in Experimental Economics* (JAI Press).
- Bossaerts, Peter, and William R. Zame, 2006, Asset trading volume in infinite-horizon economies with dynamically complete markets and heterogeneous agents: Comment, *Finance Research Letters* 3, 96 – 101.
- Bruguier, Antoine J., Steven R. Quartz, and Peter Bossaerts, 2010, Exploring the nature of “trader intuition”, *The Journal of Finance* 65, 1703–1723.
- Camerer, Colin F., and Keith Weigelt, 1996, *Research in Experimental Economics*, chapter An asset market test of a mechanism for inducing stochastic horizons in experiments, 213–238 (JAI Press).
- Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Cochrane, John H., 2001, *Asset Pricing* (Princeton University Press).
- Crockett, Sean, and John Duffy, 2013, An experimental test of the lucas asset pricing model, *Working Paper Series* .

- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *The Journal of Finance* 53, 1839–1885.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *The Journal of Finance* 45, 379–395.
- Duffie, Darrell, and Chi-Fu Huang, 1985, Implementing arrow-debreu equilibria by continuous trading of few long-lived securities, *Econometrica* 53, 1337–1356.
- Dunn, Kenneth B., and Kenneth J. Singleton, 1986, Modeling the term structure of interest rates under non-separable utility and durability of goods, *Journal of Financial Economics* 17, 27 – 55.
- Epstein, Larry G., and Tan Wang, 1994, Intertemporal asset pricing under knightian uncertainty, *Econometrica* 62, pp. 283–322.
- Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, pp. 263–286.
- Geanakoplos, John, and William R. Zame, 2007, Collateralized asset markets, Technical report, UCLA.
- Goyal, Amit, and Ivo Welch, 2003, Predicting the equity premium with dividend ratios, *Management Science* 49, 639–654.
- Grandmont, Jean Michel, 1977, Temporary general equilibrium theory, *Econometrica* 45, pp. 535–572.
- Hansen, Lars, and K. J. Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy* 91, 249–265.
- Hassan, Tarek A., and Thomas M. Mertens, 2010, The social cost of near-rational investment: Why we should worry about volatile stock markets, Working paper.
- Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *The Journal of Finance* 55, pp. 1163–1198.
- Hernandez, Alejandro, and Manuel S. Santos, 1996, Competitive equilibria for infinite-horizon economies with incomplete markets, *Journal of Economic Theory* 71, 102 – 130.
- Holt, Charles, and Susan Laury, 2002, Risk aversion and incentive effects, *American Economic Review* 92.
- Judd, Kenneth L., Felix Kubler, and Karl Schmedders, 2003, Asset trading volume with dynamically complete markets and heterogeneous agents, *The Journal of Finance* 58, 2203–2218.

- Kehoe, Timothy J., and David K. Levine, 1993, Debt-constrained asset markets, *The Review of Economic Studies* 60, 865–888.
- LeRoy, Stephen F., and Richard D. Porter, 1981, The present-value relation: Tests based on implied variance bounds, *Econometrica* 49, 555–574.
- Levine, David K., and William R. Zame, 1996, Debt constraints and equilibrium in infinite horizon economies with incomplete markets, *Journal of Mathematical Economics* 26, 103 – 131, [jce:title;Special Issue on Equilibrium with Incomplete Markets and an Infinite State Space;/ce:title](#).
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *The Review of Financial Studies* 1, pp. 41–66.
- Lucas, Robert E. Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *The Journal of Finance* 60, pp. 1167–1219.
- Maenhout, Pascal J., 2004, Robust portfolio rules and asset pricing, *The Review of Financial Studies* 17, pp. 951–983.
- Magill, Michael, and Martine Quinzii, 1994, Infinite horizon incomplete markets, *Econometrica* 62, pp. 853–880.
- Malkiel, Burton G., 1999, *A Random Walk Down Wall Street* (W.W. Norton, New York, NY).
- Martin, Ian, 2011, The lucas orchard, Working Paper 17563, National Bureau of Economic Research.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145 – 161.
- Merton, Robert C., 1973a, An intertemporal capital asset pricing model, *Econometrica* 41, pp. 867–887.
- Merton, Robert C., 1973b, Theory of rational option pricing, *The Bell Journal of Economics and Management Science* 4, 141–183.
- Nagel, Stefan, and Kenneth J. Singleton, 2011, Estimation and evaluation of conditional asset pricing models, *The Journal of Finance* 66, 873–909.

- Payzan-Le Nestour, Elise, and Peter Bossaerts, 2014, Learning about unstable, publicly unobservable payoffs, Technical report.
- Plott, Charles, and Shyam Sunder, 1988, Rational expectations and the aggregation of diverse information in laboratory security markets, *Econometrica* 56, 1085–1118.
- Rabin, Matthew, 2000, Risk aversion and expected-utility theory: A calibration theorem, *Econometrica* 68, 1281–1292.
- Radner, Roy, 1972, Existence of equilibrium of plans, prices, and price expectations in a sequence of markets, *Econometrica* 40, 289–303.
- Routledge, Bryan R, and Stanley E. Zin, 2011, Generalized disappointment aversion and asset prices, *Journal of Finance* In press.
- Samuelson, Paul A., 1973, Proof that properly discounted present values of assets vibrate randomly, *The Bell Journal of Economics and Management Science* 4, pp. 369–374.
- Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *The American Economic Review* 71, 421–436.
- Smith, Vernon L., 1965, Experimental auction markets and the walrasian hypothesis, *The Journal of Political Economy* 73, 387–393.
- Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams, 1988, Bubbles, crashes, and endogenous expectations in experimental spot asset markets, *Econometrica* 56, 1119–1151.

Appendix

Appendix 1: Theory

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

A. 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 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.

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

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

$$p_{A,t}^\sigma = \beta E \left\{ \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} \right] (d_A + p_{A,t+1}) \right\} \quad (5)$$

for $\sigma = H, L$ and $A = B, T$. (5) 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 (5) is independent of the choice of agent i , and also that we could write (5) in terms of the utility function of a representative agent – but the utility function of the representative agent would be determined *in equilibrium*.)

C. Asset prices are stationary.

Fix an asset $A = B, T$ and a period t . The stochastic Euler equation (5) 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

$$\begin{aligned}
p_{A,t}^\sigma &= \sum_{\tau=0}^{\infty} \beta^{\tau+1} E \left[\frac{u'_i(c_{i,t+\tau+1})}{u'_i(c_{i,t}^\sigma)} d_{A,t+\tau+1} \right] \\
&= \beta E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right] \sum_{\tau=0}^{\infty} \beta^\tau \\
&= \left(\frac{\beta}{1-\beta} \right) E \left[\frac{u'_i(c_i)}{u'_i(c_i^\sigma)} d_A \right]
\end{aligned} \tag{6}$$

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

D. 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 (6) 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

$$p_A^H/p_A^L = \mu \tag{7}$$

Substituting and solving yields

$$\begin{aligned}
p_A^H &= \left(\frac{\beta}{1-\beta} \right) \left[\pi d_A^H + (1-\pi)d_A^L \mu \right] \\
p_A^L &= \left(\frac{\beta}{1-\beta} \right) \left[\pi d_A^H (1/\mu) + (1-\pi)d_A^L \right]
\end{aligned} \tag{8}$$

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

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

Only one parameter, μ , determines prices of both securities in both states. Notice: $p_T^H = 2.5$ (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).

E. Asset prices are perfectly correlated with fundamentals.

This is also an immediate consequence of equations (8); 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).

F. The Tree is cheaper than the Bond.

This too is a consequence of equations (8). 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.

G. Expected returns on the Tree and Bond change with the state; they are countercyclical.

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

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

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

H. The equity premium is positive and counter-cyclical.

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).³⁴ The conclusion that the Tree is cheaper than the Bond implies that the equity premium is positive. Because asset prices are stationary, equity premia are stationary as well; simple computations show that the equity premia in the High and Low states are (remember that the expected dividends are the same for both assets and equal to d_B , and that for each asset A , $p_A^H = \mu p_A^L$):

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

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

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

This difference is strictly negative (because $\mu > 1$) so the equity premium is 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.

I. 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 probabili-

³⁴Mehra and Prescott (1985) use a slightly different model, with long-lived Tree and a one-period bond, and define the equity premium as the difference between the expected return on the risky security and that of the one-period bond.

ties, 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.³⁵

J. Agents smooth consumption over time.

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

K. Agents trade to hedge fundamental risk.

If prices were constant, agents could smooth consumption simply by buying or selling one asset. However, there *is* price risk, because prices move with fundamentals and fundamentals are uncertain. Hence, when agents sell assets because private income is low (relative to average private income), they also need to insure against the risk that prices might change by the time they are ready to buy back the assets. In equilibrium, prices increase with the dividend on the Tree, and agents correctly anticipate this. Since the Tree pays a dividend when prices are high, it is the perfect asset to hedge price risk. Consequently – but perhaps counter-intuitively! – agents *buy* Trees in periods when private income is low and *sell* when private income is high.

Appendix 2: Numerical Example

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

- There are an even number $n = 2m$ of agents; agents $i = 1, \dots, m$ are of Type I, 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 pur-

³⁵To 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.

Table XII
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%

poses. 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 XII provides equilibrium asset prices, the discounts in the price of the Tree relative to the Bond, and equity premia, as functions of the state and of risk aversion. As expected, Trees are always cheaper than Bonds. The discount on the Tree is higher in state H than in state L , while the equity premium is lower in state H than in state L , reflecting the procyclical behavior of the discount and the counter-cyclical behavior of the equity premium. The dependence of prices on the state, and the predictability of returns is apparent from the table.³⁶

Table XIII 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.^{37,38}

³⁶From Equation 5, 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.

³⁷Notice that equilibrium holdings and trade depend on whether the period is odd/even but *not* on the state (dividend of the Tree).

³⁸In 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 XIII

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.

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.

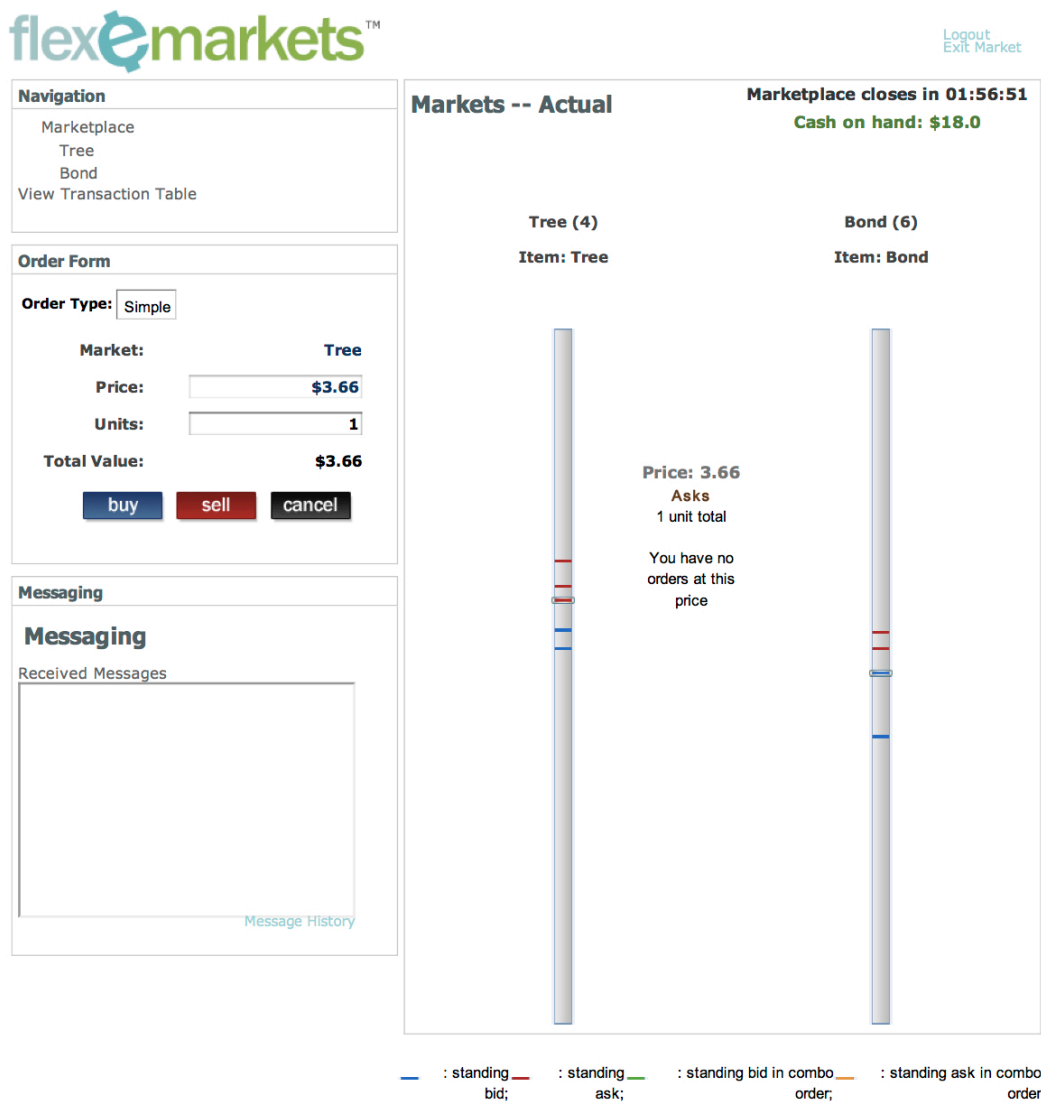


Figure 2. Time series of transaction prices of the Tree (solid line) and the Bond (dashed line); averages per period. Session numbers underneath line segments refer to Table I.




Figure 3. Time series of Tree (blue) and Bond (green) 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, in red, is the evolution of the state (High = 1; Low = 0).

