

Experiments With The Lucas Asset Pricing Model

Elena Asparouhova* Peter Bossaerts† Nilanjan Roy‡
William Zame§

December 5, 2011

Abstract

For over thirty years, the model of Lucas (1978) has been the platform of research on dynamic asset pricing and business cycles. This model restricts the intertemporal behavior of asset prices and ties those restrictions to cross-sectional behavior (the “equity premium”). The intertemporal restrictions reject the strictest interpretation of the Efficient Markets Hypothesis, namely, that prices should follow a martingale. Instead, prices move with economic fundamentals, and to the extent that these fundamentals are predictable, prices should be too. The Lucas model also prescribes the investment choices that facilitate smoothing of consumption over time. Here, we report results from experiments designed to test the primitives of the model. Our design overcomes, in novel ways, challenges to generate demand for consumption smoothing in the lab, and to induce stationarity in spite of the finite duration of lab experiments. The experiments provide overall support for the price and allocation predictions of the model, although the intertemporal price predictability is too small given observed cross-sectional price differences. Investment choices are, in turn, consistent with the low level of price predictability.¹

*University of Utah

†Caltech and SFI at EPFL

‡Caltech

§UCLA

¹Financial support from Inquire Europe, the Hacker Chair at the California Institute of Technology (Caltech), and the Development Fund of the David Eccles School of Business at the University of Utah is gratefully acknowledged.

1 Motivation

Over the last thirty years, the Lucas model (Lucas, 1978) has been the main platform that has guided the empirical research on dynamic asset pricing and business cycles. It has also become the dominant source of inspiration to financial regulators and central bankers for policy formulation.

The Lucas model delivers the core cross-sectional prediction of virtually all static asset pricing models, namely that a stock's expected return increases in the covariation ("beta") of this return with the aggregate consumption. Importantly, however, the Lucas model also provides an extension to the static approach, in making clear predictions about the intertemporal behavior of asset prices, and linking those to the cross-sectional restrictions. Specifically, it predicts that prices should co-move with economic fundamentals (aggregate consumption), and the amount of co-movement should increase with risk aversion. As such, if cross-sectional dispersion in expected returns is high because risk aversion is high, then the time-series co-movement between prices and economic fundamentals should be high as well. An immediate consequence is that prices will become predictable from the moment economic fundamentals are predictable.

The latter insight is what makes the Lucas model an invaluable formal framework within which to gauge the true empirical content of the Efficient Markets Hypothesis (EMH; Fama (1991)). Contrary to early versions of the EMH, prices need not follow a random walk (Malkiel, 1999) or even form a martingale (Samuelson, 1973). As Lucas criticizes in Section 8 of his article (Lucas, 1978): "Within this framework, it is clear that the presence of a diminishing marginal rate of substitution [...] is inconsistent with the [martingale] property."

The Lucas model is the equilibrium outcome from exchange between investors who each solve a complex dynamic programming problem whereby consumption is smoothed as much as possible given available securities, income flows, knowledge of the nature of dividend and income processes, and correct anticipation of (equilibrium) price processes. The latter makes the Lucas equilibrium an instantiation of Radner's perfect foresight equilibrium (Radner, 1972). Because they result from the solutions of complex problems, equilibrium investment policies can be very complicated, and may exhibit the *hedging* features that were first identified to be important to dynamic asset pricing in Merton (1973b) and that have become the core of the modern theory of derivatives analysis (Black and Scholes, 1973; Merton, 1973a).

On the empirical side, tests of the Lucas model have invariably been applied to historical price (and consumption) series in the field. Little attention has been paid to

its choice (investment) predictions. Starting with Mehra and Prescott (1985), the fit has generally been considered to be poor. Attempts to “fix” the model have concentrated on the auxiliary assumptions rather than on its primitives. Some authors have altered the original preference specification (time-separable expected utility) to allow for, among others, time-nonseparable utility (Epstein and Zin, 1991), loss aversion (Barberis et al., 2001), or utility functions that assign an explicit role to an important component of human behavior, namely, emotions (such as disappointment; Routledge and Zin (2011)). Others have looked at measurement problems, extending the scope of aggregate consumption series in the early empirical analysis (Hansen and Singleton, 1983), to include nondurable goods (Dunn and Singleton, 1986), or acknowledging the dual role of certain goods as providing consumption as well as collateral services (Lustig and Nieuwerburgh, 2005). Included in this category of “fixes” should be the long-run risks model of (Bansal and Yaron, 2004) because it is based on difficulty in recovering an alleged low-frequency component in consumption (growth).

Evidently, this body of empirical and theoretical research does not question the primitives of the model, namely, the claim that markets settle on a stationary Radner equilibrium where prices are measurable functions of fundamentals. The veracity of equilibration would be difficult to test on field data anyway. The problem is that we cannot observe the structural information (aggregate supply, beliefs about dividend processes, etc.) that is crucial to knowing whether markets settle at an equilibrium. (This is closely related to the Roll critique (Roll, 1977).) By contrast, laboratory experiments provide control over and knowledge of all important variables.

Thus, the first goal of the research reported on here is to bring the Lucas model to the laboratory and to test its primitives.

Since Keim and Stambaugh (1986), empirical research on the time-series properties of asset prices has been confirming that returns *are* predictable – or at least have a significant predictable component, even if such predictability is elusive at times because it is hard to recover out-of-sample (Bossaerts and Hillion, 1999). Two possible explanations have been advanced for this predictability. One points to the Lucas model and argues, like we did above, that predictability is implied by equilibrium (Bossaerts and Green, 1989; Berk and Green, 2004). Indeed, even some of the most puzzling aspects of predictability can be shown to be equilibrium implications of simple assumptions about the structure of the economy (Brav and Heaton, 2002; Li et al., 2009). The second one is that predictability is an aggregate expression of the many cognitive biases that have been demonstrated at the individual level; this is the core thesis of Behavioral Finance (Bondt and Thaler, 1985).

The second goal of our research, then, is to inform the controversy about predictability of returns in asset markets as it relates to the EMH. Can we generate predictability in the laboratory, and if so, is its nature consistent with the Lucas model?

The design of such an experiment is challenging, because the Lucas model assumes that the world is stationary, continues forever, and that investment demands are driven primarily by the desire to smooth consumption. Of these challenges, the infinite horizon is the easiest to deal with: one merely has to introduce a stochastic ending time; see Camerer and Weigelt (1996). The finite experiment duration, however, makes stationarity particularly difficult to induce, as beliefs necessarily change when time approaches the officially announced termination of the experiment. Likewise, it is difficult to imagine that participants care when they receive their earnings over the course of the experiment, which potentially negates the assumption of preferences for earnings smoothing.

We introduce novel features to the design of an intertemporal asset pricing experiment that overcome these challenges. Their validity hinges in part on an important component of the (original) Lucas model, namely, time-separable utility. We make sure that time separability follows naturally from our design, and would only fail if participants did not have expected utility preferences.

The proposed experiment is related to Crockett and Duffy (2010). There are at least two major differences, however. First, we do not induce demand for smoothing by means of nonlinearities in earnings as a function of experiment payoffs, but generate it as the result of a novel experimental design feature. The predicted cross-sectional pricing patterns are therefore driven entirely by the uncertainty of the dividends of (one of) the assets, exactly as in the original Lucas model, unconfounded by nonlinearities. Second, we explicitly address the issue of finite experiment duration when implementing a stationary termination protocol.

The remainder of this paper is organized as follows. Section 2 introduces the Lucas model by means of a stylized example that will form the basis of the experiment. Section 3 details the experimental design. Section 4 presents the results from a series of six experimental sessions. Section 5 provides discussion. The last section concludes.

2 The Lucas Asset Pricing Model

We envisage an environment with minimal complexity yet one that generates a rich set of predictions about prices across time and allocations across types of investors. Perhaps most importantly, the environment is such that trading is necessary in *each*

period. Inspired by Bossaerts and Zame (2006), we want to avoid a situation (as in Judd et al. (2003)), where theory predicts that trade will take place only once. When bringing the setting to the laboratory, it would indeed be rather awkward to give subjects the opportunity to trade every period while the theory predicts that they should not!

For our theoretical benchmark, we consider a stationary, infinite horizon economy in which infinite-lived agents with time-separable expected utility are initially allocated two types of assets: (i) a *Tree* that pays a stochastic dividend of \$1 or \$0 every period, each with 50% chance, independent of past outcomes, and (ii) a (consol) *Bond* that always pays \$0.50.

There is an *equal number* of two types of agents. Type I agents receive income of \$15 in even periods (2, 4, 6,...), while those of Type II receive income of \$15 in odd periods. As such, total (economy-wide) income is constant over time. Before period 1, Type I agents are endowed with 10 Trees and no Bonds; Type II agents start with 0 Trees and 10 Bonds.

Assets pay dividends $d_{k,t}$ ($k \in \{\text{Tree}, \text{Bond}\}$) *before* period t ($t = 1, 2, \dots$) starts. At that point, agents also receive their income, $y_{i,t}$ ($i = 1, \dots, I$), as prescribed above. As dividends and income are fungible, we refer to them as *cash*, and cash is perishable. In what follows, $c_{i,t}$ denotes the cash available to agent i in period t . Agents have *common* time-separable utility for cash:

$$U_i(\{c_{i,t}\}_{t=1}^{\infty}) = E \left\{ \sum_{t=1}^{\infty} \beta^{t-1} u(c_{i,t}) \right\}. \quad (1)$$

Markets open and agents can trade their Trees and Bonds for cash, subject to a standard budget constraint. To determine optimal trades, agents take asset prices $p_{k,t}$ ($k \in \{\text{Tree}, \text{Bond}\}$) as given, and correctly anticipate (à la Radner (1972)) that future prices are a time-invariant function of the only variable economic fundamental in the economy, namely, the dividend on the Tree $d_{\text{Tree},t}$. In particular they know that prices are set as follows:

$$p_{k,t} = \beta E \left[\frac{u'(c_{i,t+1})}{u'(c_{i,t})} (d_{k,t+1} + p_{k,t+1}) \right]. \quad (2)$$

We shall not go into details here, because the derivation of the equilibrium is standard. Instead, here are the main predictions of the resulting (Lucas) equilibrium. For the parametric illustrations, we set $\beta = 5/6$, and we assume constant relative risk aversion; if risk aversion equals 1, agents are endowed with logarithmic utility ($u(c_{i,t}) = \log(c_{i,t})$).

Table 1: Equilibrium Prices And ‘Equity Premium’ As A Function Of (Constant Relative) Risk Aversion And State (Level Of Dividend On Tree).

Risk Aversion	State	Tree	Bond	(‘Equity Premium’)
0.1	High	2.50	2.55	(0.05)
	Low	2.40	2.45	(0.05)
	(Difference)	(0.10)	(0.10)	(0)
0.5	High	2.50	2.78	(0.22)
	Low	2.05	2.27	(0.22)
	(Difference)	(0.45)	(0.51)	(0)
1	High	2.50	3.12	(0.62)
	Low	1.67	2.09	(0.42)
	(Difference)	(0.83)	(1.03)	(0.20)

1. **Cross-sectional Restrictions:** Because the return on the Tree has higher co-variability (or “beta”) with aggregate consumption (which varies only because of the dividend on the Tree), its equilibrium price is lower than that of the Bond, replicating a well-known result from static asset pricing theory. Note that this result is far from trivial: returns are determined not only by future dividends, but also future *prices*, and it is not *a priori* clear that prices behave like dividends! With logarithmic utility, the difference between the price of the Tree and that of the Bond is \$0.62 if the dividend on the Tree is high (\$1), and \$0.42 when this dividend is low (\$0). See Table 1. This table also lists prices and corresponding ‘equity premia’ for risk aversion coefficients equal to 0.5 (square-root utility) and 0.1.² To be somewhat consistent with prior work, we refer to the difference between the Bond and Tree prices as the ‘equity premium.’ Usually, the equity premium is defined as the difference in the expected yield (return) on risky securities and the riskfree rate. Because of the inverse relationship between yield and price, we here refer to the ‘equity premium’ as the difference between the Bond price and the Tree price. Our ‘equity premium’ increases in the standard equity premium.

2. **Intertemporal Restrictions:** Asset prices depend on the dividend on the Tree.

²The equilibrium prices are unique; in particular, they do not depend on the State outcome in Period 1 (State = dividend on tree).

With logarithmic utility, the Tree price is \$2.50 when the Tree dividend is high, and \$1.67 when it is low; the corresponding Bond prices are \$3.12 and \$2.09. See Table 1. Such prices induce substantial predictability in the asset return series: when the dividend of the Tree is high, the expected return on the Tree is only 3.4% (equal to $(0.5 \cdot (2.50 + 1) + 0.5 \cdot 1.67) / 2.5 - 1$) while it equals 55% when the dividend on the Tree is low! This predictability contrasts with simple formulations of EMH (Fama, 1991) which posit that expected returns are constant.

3. **Linking Cross-sectional and Intertemporal Restrictions:** as risk tolerance increases, the (cross-sectional) difference between the prices of the Tree and the Bond diminishes, as does the time-series dependence of prices on economic fundamentals. Table 1 shows how the difference in prices of an asset decreases with risk aversion (the Tree price difference decreases from 0.83 to 0.45 and 0.10 as one moves from logarithmic utility down to risk aversion equal to 0.5 and 0.1) while at the same time the ‘equity premium’ (averaged across states) drops from 0.52 to 0.22 to 0.05. In the extreme case of risk neutrality, both the Tree and Bond are priced at a constant \$2.50. As a matter of fact, for the range of risk aversion coefficients between 0 (risk neutrality) and 1 (logarithmic utility), the correlation between the difference in prices across states and the ‘equity premium’ (averaged across states) equals 0.99 for the Tree and 1.00 for the Bond!³
4. **Equilibrium Consumption:** In equilibrium, consumption across types should be perfectly rank-correlated. With only two (dividend) states, this means that consumption for both types is high in the high state and low in the low state. Since we assume that agents have identical preferences, they should consume a constant fraction of the aggregate cash flow (the total of dividends and incomes). Thus, agents fully offset their income fluctuations and as a result obtain smooth consumption.
5. **Trading for Consumption Smoothing:** Agents obtain equilibrium consumption smoothing mostly through exploiting the price differential between Trees and Bonds: when they receive no income, they sell Bonds and buy Trees, and since the Tree is always cheaper, they generate cash; conversely, in periods when they do receive income, they buy (back) Bonds and sell Trees, depleting their cash because Bonds are more expensive. See Table 2.⁴

³The relationship is slightly nonlinear, which explains why the correlation is not a perfect 1.

⁴Equilibrium holdings and trade do *not* depend on the state (dividend of the Tree). However, they do depend on the state in Period 1. Here, we assume that the state in Period 1 is high (i.e., the Tree pays a dividend of \$1). When the state in Period 1 is low, there is a technical problem for risk aversion of 0.5 or

Table 2: Type I Agent Equilibrium Holdings and Trading As A Function Of (Constant Relative) Risk Aversion And Period (Odd; Even).

Risk Aversion	Period	Tree	Bond	(Total)
0.1	Odd	5.17	2.97	(8.14)
	Even	4.63	6.23	(10.86)
	(Trade in Odd)	(+0.50)	(-3.26)	
0.5	Odd	6.32	1.96	(8.28)
	Even	3.48	7.24	(10.72)
	(Trade in Odd)	(+2.84)	(-5.28)	
1	Odd	7.57	0.62	(8.19)
	Even	2.03	7.78	(9.81)
	(Trade in Odd)	(+5.54)	(-7.16)	

6. **Trading to Hedge Price Risk:** Because prices move with economic fundamentals, and economic fundamentals are risky (because the dividend on the Tree is), there is price risk. When they sell assets to cover an income shortfall, agents 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 maybe counter-intuitively!), agents *buy* Trees in periods with income shortfall and they *sell* when their income is high. See Table 2, which shows, for instance, that a Type I agent with logarithmic preferences will *purchase* more than 5 Trees in periods when they have no income (Odd periods), subsequently selling them (in Even periods) in order to buy back Bonds. Hedging is usually associated with Merton's intertemporal asset pricing model and is the core of modern derivatives analysis. Here, it forms an integral part of the trading predictions in the Lucas model.

In the experiment, we tested the above six, inter-related predictions.

higher: in Odd periods, agents need to short sell Bonds. In the experiment, short sales were not allowed.

3 Implementing the Lucas Model

When planning to implement the above Lucas economy, three issues emerge.

- a. Because actual consumption is not feasible until after the session concludes, it would not make much of a difference if we were to pay subjects' earnings gradually. So, there is no natural demand for consumption smoothing in a laboratory experiment. This is why Crockett and Duffy resort to nonlinearities in payoff-earnings relationships (Crockett & Duffy, 2010). Ideally, however, one would like to avoid this, because nonlinearities are not part of the original Lucas model.
- b. The Lucas economy has an infinite horizon, but an experimental session has to end in finite time. A relatively simple fix exists for this problem, which is to end the session stochastically (Camerer and Weigelt, 1996). This ending procedure also introduces discounting: the discount factor will be proportional to the probability of continuing the session.
- c. The Lucas economy is stationary. In principle, this is also easy to fix, by applying a constant termination probability: the chance that a period is the terminal one for the session is constant; in particular, it does not depend on how long the session has been lasting. But this is not satisfactory, because it means that the experiment could go on forever, or at least, take much longer than a typical experimental session (2-3 hours).

In our experiment, we used the standard solution to resolve issue (b), which is to randomly determine if a period is terminal. We set the termination probability equal to $1/6$, which means that we induced a discount factor of $\beta = 5/6$ (the number used in the theoretical calculations in the previous section). In particular, after the markets in period t close, 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 (a), we made end-of-period individual cash holdings *perish* in each period that was not terminal; only securities holdings carried over to the next period. If a period *was* terminal, however, securities holdings perished. Participants' earnings were then determined entirely by the cash they held at the end of this terminal period. As such, if participants have expected utility preferences, their preferences will automatically become of the time-separable type that Lucas uses in his model, albeit with an adjusted discount factor: the period- t discount factor becomes $(1 - \beta)\beta^{t-1}$. It is straightforward to show that all results (prices; allocations) remain the same, simply because the new utility function to be maximized is proportional to the old one

[Eqn. (1)] with constant of proportionality $(1 - \beta)$.

It is far less obvious how to resolve problem (c). Here, we propose a simple solution, exploiting essential features of the Lucas model. It works as follows. We announce that the experimental session will last until a pre-specified time and there will be as many replications of the (Lucas) economy as can be fit within this time. If a replication finishes at least 10 minutes before the announced end time, a new replication starts. Otherwise, the experimental session is over. *If a replication is still running by the closing time, we announce before trade starts that the current period is either the last one (if our die shows 7 or 8) or the penultimate one (for all other values of the die).* In the latter case, we move to the next period and this one becomes the terminal one with certainty. This means that subjects will keep the cash they receive through dividends and income for that period. (There will be no trade because assets perish at the end.)

It is straightforward to show that the equilibrium prices remain the same whether the new termination protocol is applied or if termination continues to be determined each period with the roll of a die. In the former case, the pricing formula is:⁵

$$p_{k,t} = \frac{\beta}{1 - \beta} E\left[\frac{u'(c_{i,t+1})}{u'(c_{i,t})} d_{k,t+1}\right]. \quad (3)$$

To see that the above is the same as the formula in Eqn. (2), apply the assumption of i.i.d. dividends and the consequent stationary investment rules (which generate i.i.d. consumption flows) to re-write Eqn. (2) as follows:

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

which is the same as Eqn. (3).

⁵To derive the formula, consider agent i 's optimization problem in period t , which is terminal with probability $1 - \beta$, and penultimate with probability β , namely: $\max (1 - \beta)u(c_{i,t}) + \beta E[u(c_{i,t+1})]$, subject to a standard budget constraint. The first-order conditions are, for asset k :

$$(1 - \beta) \frac{\partial u(c_{i,t})}{\partial c} p_{k,t} = \beta E\left[\frac{\partial u(c_{i,t+1})}{\partial c} d_{k,t+1}\right].$$

The left-hand side captures expected marginal utility from keeping cash worth one unit of the security; the right-hand side captures expected marginal utility from buying the unit; for optimality, the two expected marginal utilities have to be the same. Formula (3) obtains by re-arrangement of the above equation. Under risk neutrality, and with $\beta = 5/6$, $p_{k,t} = 2.5$ for $k \in \{\text{Tree, Bond}\}$

Table 3: Summary data, all experimental sessions.

Session	Place	Replication Number	Periods (Total, Min, Max)	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 ran as many replications as possible within the time allotted to the experimental session. In order to avoid wealth effects on subject preferences, we paid for only a fixed number (say, 2) of the replications, randomly chosen after conclusion of the experiment. (If we ran less replications than this fixed number, we paid multiples of some or all of the replications.)

Sample instructions are included in this paper as an appendix.

4 Results

Until now (August 2011), we have conducted six experimental sessions, with the participant number ranging between 12 and 30. Three sessions were conducted at Caltech, two at UCLA, and one at the University of Utah. This generated 80 periods in total, spread over 15 replications. Table 3 provides specifics. Our novel termination protocol was applied in all sessions. The starred sessions ended with a period in which participants knew for sure that it was the last one, and hence, generated no trade.

We now present the experimental results.

Volume. Table 4 lists average trading volume per period (excluding periods in which should be no trade). Consistent with theoretical predictions, trading volume in Periods 1 and 2 is significantly higher; it reflects trading needed for agents to move to their steady-state holdings. In the theory, subsequent trade takes place only to smooth consumption across odd and even periods. Volume in the Bond is significantly lower in Periods 1 and 2; this is an artefact of the few replications when the state in Period 1 was low. This deprived Type I participants of cash (Type I participants start with

Table 4: Trading volume.

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

10 Trees and no income). In principle, they should have been able to sell enough Trees to buy Bonds, but evidently they did not manage to complete all the necessary trades in the allotted time (four minutes).

Cross-Sectional Price Differences. Table 5 displays average period transaction prices as well as the period’s state (“high” if the dividend of the Tree was \$1; “low” if it was \$0). The table demonstrates that, consistent with the Lucas model, the Bond is priced *above* the Tree (see Prediction 1 above), with price differential (‘equity premium’) of about \$0.50. When checking against Table 1, this reflects a (constant relative) risk aversion coefficient of 1.

Prices Over Time. Figure 1 shows a plot of the evolution of (average) prices over time, arranged chronologically by experimental sessions (numbered as in Table 3); replications within a session are concatenated. The plot reveals that prices are noisy. 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 (Prediction 2). Specifically,

Table 5: Period-average transaction prices and corresponding ‘equity premium’.

	Tree Price	Bond Price	‘Equity Premium’
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

Table 6: Mean period-average transaction prices and corresponding ‘equity premium,’ as a function of state.

State	Tree Price	Bond Price	‘Equity Premium’
High	2.91	3.34	0.43
Low	2.66	3.20	0.54
Difference	0.24	0.14	-0.11


prices should be high in high states, and low in low states. In reality, much more is going on. In particular, contrary to the Lucas model, we detect price drift. The direction of the drift is not always obvious; the drift appears to be stochastic.

Nevertheless, behind the noise, clear evidence in favor of the Lucas model emerges. As Table 6 shows, prices in the high state are on average 0.24 (Tree) and 0.14 (Bond) above those in the low state. The table does not display statistical information; (average) transaction prices are not i.i.d., so that we cannot rely on standard t tests to determine significance. We will provide formal statistical evidence later on.

While prices in ‘good’ periods are higher than in ‘bad’ ones, the differential is too small compared to the size of the ‘equity premium.’ The average ‘equity’ premium of \$0.50 reveals a coefficient of relative risk aversion of 1, and this level of risk aversion would imply a price differential across states of \$0.83 and \$1.03 for the Tree and Bond, respectively. See Table 1. In the data, the price differentials amount to only \$0.24 and \$0.14. In other words, while there is predictability in price series, the co-movement between prices and economic fundamentals is far lower than the Lucas model predicts.

Both Cross-Sectional And Time Series Price Properties. Still, the theory

Figure 1: Time series of Tree (solid line) and Bond (dashed line) transaction prices; averages per period. Session numbers underneath line segments refer to Table 3.



also states that the differential in prices between ‘good’ and ‘bad’ periods should increase with the ‘equity premium’ (Prediction 3). Table 7 shows that this is true in the data. The observed correlation is not perfect (unlike in the theory), but marginally significant for the Tree; it is insignificant for the Bond.

Table 6 also shows that the ‘equity premium’ is higher in periods when the state is low than when it is high. This is inconsistent with the theory. The average level of the ‘equity premium’ reveals logarithmic utility, and for this type of preferences, the equity premium should be *lower* in bad periods; see Table 1.⁶

Prices: Formal Statistics. To enable formal statistical statements about the

⁶The requirement that the equity premium is higher in ‘good’ periods does not depend on the level of risk aversion. But for lower levels of risk aversion, the difference in ‘equity premium’ is hardly detectible, as is evident from Table 1.

Table 7: Correlation between equity premium (average across periods) and price differential of tree and bond across high and low periods.

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

price differences in ‘good’ and ‘bad’ states, we ran a regression of period transaction price levels onto the state (=1 if high; 0 if low). To adjust for the time series dependence revealed in Figure 1, we add session dummies and a time trend (Period number). In addition, to gauge the effect of our session termination protocol, we add a dummy for periods when we announce that the session is about to come to a close, and hence, the period is either the penultimate or last one, depending on the draw of the die. Lastly, we add a dummy for even periods. Table 8 displays the results.

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

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

Closer inspection of the properties of the error term did reveal substantial dependence over time, despite our including dummies to mitigate time series effects. Table 8 shows Durbin-Watson (DW) test statistics with value that correspond to $p < 0.001$.

Proper time series specification analysis revealed that the best model involved first differencing price changes, effectively confirming the stochastic drift evident in Figure 1. All dummies could be deleted, and the highest R^2 was obtained when explaining (average) price changes as the result of a *change* in the state. See Table 9.⁷ For the Tree, the effect of a change in state from low to high is a significant \$0.19 ($p < 0.05$). The effect of a change in state on the Bond price remains insignificant, however ($p > 0.05$). The autocorrelations of the error terms are now acceptable (marginally above their standard errors).

⁷We deleted observations that straddled two replications. Hence, the results in Table 9 are solely based on intra-replication price behavior. The regression does not include a dummy; average price changes are insignificantly different from zero.


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

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

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

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

Figure 2: Time series of period-average Tree (red line) and Bond (blue line) transaction price changes. Changes are concatenated across all replications and all sessions, but exclude inter-replication observations. State is indicated by black solid line on top; state = 2 when “high” (tree dividend equals \$1); state = 0 when “low” (tree dividend equals \$0).



At 18%, the explained variance of Tree price changes (R^2) is high. In theory, one should be able to explain 100% of price variability. But prices are noisy, as Figure 1 revealed. The regression in first differences shows that there fundamental economic forces are behind price changes, at least as far as the Tree is concerned, as predicted by the Lucas model.

Figure 2 displays the evolution of price changes, after chronologically concatenating all replications for all sessions. Like in the data underlying the regression in Table 9, the plot only shows intra-replication price changes. The period state (=1 if low; 2 if high) is plotted on top.

Consumption Across States. Prediction 4 of the Lucas model states that agents

Table 10: Average consumption (end-of-period cash holdings) as a function of participant Type and State. Autarky numbers in parentheses.

Type	Consumption (\$)		Consumption Ratio	
	High	Low	High	Low
I	14.93 (19.75)	7.64 (4.69)	1.01 (0.52)	1.62 (3.26)
II	15.07 (10.25)	12.36 (15.31)		

of both types should trade to holdings that generate high consumption in high states, and low consumption in low states. Assuming identical preferences, they should consume a fixed fraction of total period cash flows, or the ratio of Type I to Type II consumption should be equal in both states. The left-hand panel of Table 10 displays the average amount of cash per type in high vs. low states.⁸ In parentheses, we indicate consumption levels assuming that agents do not trade (i.e., under an autarky). The statistics in the table confirm that consumption of *both* types increases with dividend levels. The result is economically significant because consumption is *anti*-correlated under autarky. This is strong evidence in favor of Lucas' model.

The right panel of Table 10 displays Type II's consumption as a ratio of Type I's consumption. The difference is substantially reduced from what would obtain under autarky (which is displayed in parentheses). Again, this supports the Lucas model, though the theory would want the consumption ratios to be exactly equal across states.

There are significant individual differences however, reminiscent of the huge cross-sectional variation in choices in static asset pricing experiments that led to the development of the ε -CAPM (Bossaerts et al., 2007).

Consumption Across Odd And Even Periods. Prediction 4 states that subjects should be able to perfectly offset income differences across odd and even periods. Table 11 demonstrates that our subjects indeed managed to smooth consumption substantially; the outcomes are far more balanced than under autarky (in parentheses; averaged across high and low states, excluding Periods 1 and 2).

Price Hedging. The above results suggest that our subjects (on average) managed to move substantially towards equilibrium consumption patterns in the Lucas model. However, they did not resort to price hedging (Prediction 6) as a means to ensure those patterns. Table 12 displays average asset holdings across periods for Type I

⁸To compute these averages, we ignored Periods 1 and 2, to allow subjects time to trade from their initial holdings to steady state positions.

Table 11: Average consumption (end-of-period cash holdings) as a function of participant and period Types.

Type	Consumption (\$)	
	Odd	Even
I	7.69 (2.41)	13.91 (20.65)
II	14.72 (20)	11.74 (5)

Table 12: Average (End-Of-Period) Asset Holdings Of Type I Participants.

Period	1	2	3	4	5	6	7	8	9
Income (\$)	0	15	0	15	0	15	0	15	0
Asset (Initial Holding):									
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

subjects (who receive income in even periods). They are net sellers of assets in periods of income shortfall (see “Total” row), just like the theoretical agents with logarithmic utility (see Table 2). But unlike in the theory, subjects *decrease* Tree holdings in low-income periods and *increase* them in high-income periods (compare to Table 2). As a by-product, Type I subjects generate cash mostly through selling Trees as opposed to exploiting the price differential between the Bond and the Tree.

Altogether, it appears that the findings from our experiments are in line with the predictions of the Lucas model, with two exceptions: (i) the co-movement between prices and economic fundamentals was too small; (ii) subjects did not engage in price hedging.

These two anomalies should be related. Subjects may not expect prices to change with economic fundamentals, and hence, they rationally perceive no need to hedge price risk. These beliefs are generally not falsified because of the short duration of a typical replication, despite the fact that the beliefs are actually wrong, at least as far as the Tree is concerned – but this we ourselves discovered only after pooling the price behavior from all sessions.

5 Discussion

The experiments demonstrate that our design “delivers” in the sense that it generates meaningful results in line with the Lucas model – with the exception of the much reduced co-movement of prices with economic fundamentals (given the significant differences of prices cross-sectionally) and, consistent with it, the absence of concern for price hedging among participants.

Here, we discuss potential causes of these two anomalies.

Starting with Epstein and Zin (1991), it has become standard in research on the Lucas model with historical field data to use time-nonseparable preferences. We refrain from following this tradition because in the context of our experimental design, time separability is a natural consequence of expected utility. To justify time nonseparable preferences, one would have to reject expected utility altogether. In field research, time non-separability was introduced not because researchers gave up on expected utility, but because they wanted a convenient framework in which to disentangle risk aversion and consumption smoothing.

Instead, we think that it would be more fruitful to start studying the effect of small deviations from perfect foresight. The Lucas model in general, and price hedging in particular, hinge on perfect foresight (of the equilibrium relationship between prices and dividends). One wonders what the effect is of small deviations from perfect foresight on equilibrium prices and allocations.

The absence of demand for (price) hedging may indeed reflect subjects’ belief that prices do not co-move with economic fundamentals; essentially, in accordance with the version of the EMH that the Lucas model discredits, prices are expected to be a martingale. The belief is wrong, as we pointed out, but not readily falsifiable within the short time of an experiment, and even after 80 observations (80 periods), not falsifiable for Bond prices. That is, the belief is a good working hypothesis.

We can start from an extreme. Imagine that agents always expect prices to stay at the current level, irrespective of future economic fundamentals. Agents correctly solve their dynamic investment-consumption problem (given their beliefs), send corresponding demands to markets, where prices are such that there is equilibrium each period. How would this equilibrium evolve over time? Simulations which we performed so far have revealed that prices will still co-move with dividends, albeit in a much reduced way.

The equilibrium we sketched above is obviously not a Radner equilibrium, because expectations of the mapping from states (dividend levels) to prices will be incorrect.

An interesting question is whether we can transform it into a proper Radner equilibrium through the introduction of noise: agents posit that prices do not only depend on dividends, but also on some exogenous “noise.” Can we identify types of noise which generate Radner equilibria with significant equity premia but minimal time series dependence, and hence, minimal demand for hedging?

One could envisage that the proposed “noise” reflects doubts in the mind of the agents about the veracity of the posited mapping between states and prices, in analogy with the “noise” added to conjectures of strategic behavior of other players in the quantal response equilibrium (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998).

We leave this possibility for future research.

6 Conclusion

Over the last thirty years, the Lucas model has become *the* core theoretical model through which scholars of macroeconomics and finance view the real world, advise investments in general and retirement savings in particular, prescribe economic and financial policy and induce confidence in financial markets. Despite this, little is known about the true relevance of the Lucas model. The recent turmoil in financial markets and the effects it had on the real economy has severely shaken the belief that the Lucas model has anything to say about financial markets. Calls are being made to return to pre-Lucas macroeconomics, based on reduced-form Keynesian thinking. This paper was prompted by the belief that proper understanding of whether the Lucas model (and the Neoclassical thinking underlying it) is or is not appropriate would be enormously advanced if we could see whether the model did or did not work in the laboratory.

Of course, it is a long way from the laboratory to the real world, but it should be kept in mind that no one has ever seen convincing evidence of the Lucas model “at work” – just as no one had seen convincing evidence of another key model of finance (the Capital Asset Pricing Model or CAPM) at work until the authors (and their collaborators) generated this evidence in the laboratory (Asparouhova et al., 2003; Bossaerts and Plott, 2004; Bossaerts et al., 2007). The research provides absolutely crucial – albeit modest – evidence concerning the scientific validity of the core asset pricing model underlying formal macroeconomic and financial thinking.

Specifically, despite their complexity, our experimental financial markets exhibited many features that are characteristic of the Lucas model, such as the co-existence of a significant equity premium and (albeit reduced) co-movement of prices and economic

fundamentals. Consistent with the model, the co-movement increased with the magnitude of the equity premium. And subjects managed to smooth consumption over time. Smoothing was not perfect, but sufficient for consumption to become positively correlated across subject types, in sharp contrast with consumption under autarky, which was negatively correlated. But we did not observe demand for price hedging, perhaps because subjects wrongly believed that prices were a martingale. Still, such beliefs were not irrational: within the time frame of a single session, there was insufficient evidence to the contrary.

But the overall evidence against a martingale was strong, at least for price series for the risky asset (the Tree). As such, our data reject the EMH to the extent that it requires prices to be a martingale. Prices *were* predictable. After a ‘good’ period (the dividend on the Tree is high), Tree prices were expected to decrease 3.3% on average, while after a ‘bad’ period, they were expected to increase 3.6%. This price predictability translates into strongly time-varying expected returns. After a ‘good’ period the expected return on the Tree equaled only 14%. After a ‘bad’ period, it was as high as 22%.⁹ Again, the Lucas model is consistent with this price predictability and the resulting time-variability in expected returns.

⁹These calculations are based on the results displayed in Tables 9 (time series model) and 6 (average prices). E.g., when the state is ‘high’ (i.e., the period is ‘good’), then average prices equal 2.91 (Table 6). The Tree pays a dividend on \$1 with 50% chance. The state changes to ‘low’ with 50% chance, at which point the price is expected to drop \$0.19 according to the model of Table 9. So, the expected return equals $0.50 * 1/2.91 - 0.50 * 0.19/2.91$, or 14%.

Appendix: Instructions (Type I Only)

Web Address: filagora.caltech.edu/fm/

User name:

Password:

INSTRUCTIONS

1. Situation

One session of the experiment consists of a number of replications of the same situation, referred to as *periods*.

You will be allocated *securities* that you can carry through all periods. You will also be given *cash*, but cash will not carry over from one period to another.

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

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

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

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

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

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

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

During the experiment, accounting is done in real dollars.

2. Data

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

You will start this session with **10 trees** and **0 bonds**. Others may start with different initial allocations.

In addition, you will receive income every alternate period. In **odd** periods (1,3...) you will receive nothing, and in **even** periods (2,4,...) you will receive 15 dollars. This income is added to your cash at the beginning of a new period. Others may have a different income flow.

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

3. Examples (for illustration only)

Tables 1 and 2 give two sample examples of outcomes in a session. It is assumed that the session ends after the 6th period. Table 1 shows the asset holdings, dividend and cash each period if the states are as per row 2 and the individual sticks to the initial allocation throughout. The final take-away cash/earning is 25 dollars as the session terminated after 6th period. It would have been \$0 if it had terminated in period 5.

Table 2 shows the case where the individual trades as follows:

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

Since there is a 1/6 chance that the session ends in the period when a security is bought, its expected value equals $(\frac{5/6}{1-5/6})$ times the expected dividend (which is equal for both the tree and the bond), or $(5) * (0.5) = 2.50$. Trade is assumed to take place at 2.50. Note, however, that the actual trading prices may be different, and that they may even change over time, depending on, e.g., the dividend on the tree. The final take-away cash in this case is \$15.00. It would have been \$13.00 if the session had terminated in period 5.

Table 1.

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

Table 2.

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

References

Elena Asparouhova, Peter Bossaerts, and Charles Plott. Excess demand and equilibrium in multi-security financial markets: The empirical evidence. *Journal of*

Financial Markets, 6:1–21, 2003.

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

Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116(1):1–53, 2001. ISSN 00335533. URL <http://www.jstor.org/stable/2696442>.

Jonathan B. Berk and Richard C. Green. Mutual fund flows and performance in rational markets. *Journal of Political Economy*, 112(6):1269 – 1295, 2004. ISSN 00223808.

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):pp. 637–654, 1973. ISSN 00223808. URL <http://www.jstor.org/stable/1831029>.

Werner F. M. De Bondt and Richard Thaler. Does the stock market overreact? *The Journal of Finance*, 40(3):pp. 793–805, 1985. ISSN 00221082. URL <http://www.jstor.org/stable/2327804>.

Peter Bossaerts and Richard C. Green. A general equilibrium model of changing risk premia: Theory and tests. *The Review of Financial Studies*, 2(4):467–493, 1989. ISSN 08939454. URL <http://www.jstor.org/stable/2962065>.

Peter Bossaerts and Pierre Hillion. Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies*, 12(2):405–428, 1999. doi: 10.1093/rfs/12.2.405. URL <http://rfs.oxfordjournals.org/content/12/2/405.abstract>.

Peter Bossaerts and Charles Plott. Basic principles of asset pricing theory: Evidence from large-scale experimental financial markets. *Review of Finance*, 8(2):135–169, 2004.

Peter Bossaerts and William R. Zame. Asset trading volume in infinite-horizon economies with dynamically complete markets and heterogeneous agents: Comment. *Finance Research Letters*, 3(2):96 – 101, 2006. ISSN 1544-6123. doi: DOI: 10.1016/j.frl.2006.01.001. URL <http://www.sciencedirect.com/science/article/pii/S1544612306000031>.

- Peter Bossaerts, Charles Plott, and William Zame. Prices and portfolio choices in financial markets: Theory, econometrics, experiment. *Econometrica*, 75:993–1038, 2007. Mimeo.
- Alon Brav and J.B. Heaton. Competing theories of financial anomalies. *Review of Financial Studies*, 15(2):575–606, 2002. doi: 10.1093/rfs/15.2.575. URL <http://rfs.oxfordjournals.org/content/15/2/575.abstract>.
- Colin F. Camerer and Keith Weigelt. *Research in Experimental Economics*, chapter An asset market test of a mechanism for inducing stochastic horizons in experiments, pages 213–238. JAI Press, 1996.
- Sean Crockett and John Duffy. A Dynamic General Equilibrium Approach to Asset Pricing Experiments. *SSRN eLibrary*, 2010.
- Kenneth B. Dunn and Kenneth J. Singleton. Modeling the term structure of interest rates under non-separable utility and durability of goods. *Journal of Financial Economics*, 17(1):27 – 55, 1986. ISSN 0304-405X. doi: DOI: 10.1016/0304-405X(86)90005-X. URL <http://www.sciencedirect.com/science/article/pii/0304405X8690005X>.
- Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 99(2):pp. 263–286, 1991. ISSN 00223808. URL <http://www.jstor.org/stable/2937681>.
- Eugene F. Fama. Efficient capital markets: Ii. *The Journal of Finance*, 46(5):pp. 1575–1617, 1991. ISSN 00221082. URL <http://www.jstor.org/stable/2328565>.
- Lars Hansen and K. J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91:249–265, 1983.
- Kenneth L. Judd, Felix Kubler, and Karl Schmedders. Asset trading volume with dynamically complete markets and heterogeneous agents. *The Journal of Finance*, 58(5):2203–2218, 2003. ISSN 1540-6261. doi: 10.1111/1540-6261.00602. URL <http://dx.doi.org/10.1111/1540-6261.00602>.
- Donald B. Keim and Robert F. Stambaugh. Predicting returns in the stock and bond markets. *Journal of Financial Economics*, 17(2):357 – 390, 1986. ISSN 0304-405X. doi: DOI: 10.1016/0304-405X(86)90070-X. URL <http://www.sciencedirect.com/science/article/pii/0304405X8690070X>.

Erica X. N. Li, Dmitry Livdan, and Lu Zhang. Anomalies. *Review of Financial Studies*, 22(11):4301–4334, 2009. doi: 10.1093/rfs/hhp023. URL <http://rfs.oxfordjournals.org/content/22/11/4301.abstract>.

Robert E. Jr. Lucas. Asset prices in an exchange economy. *Econometrica*, 46(6): 1429–1445, 1978. ISSN 00129682. URL <http://www.jstor.org/stable/1913837>.

Hanno N. Lustig and Stijn G. Van Nieuwerburgh. Housing collateral, consumption insurance, and risk premia: An empirical perspective. *The Journal of Finance*, 60(3):pp. 1167–1219, 2005. ISSN 00221082. URL <http://www.jstor.org/stable/3694924>.

Burton G. Malkiel. *A Random Walk Down Wall Street*. W.W. Norton, New York, NY, 1999.

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1:9–41, 1998. ISSN 1386-4157. URL <http://dx.doi.org/10.1023/A:1009905800005>. 10.1023/A:1009905800005.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6 – 38, 1995. ISSN 0899-8256. doi: DOI: 10.1006/game.1995.1023. URL <http://www.sciencedirect.com/science/article/pii/S0899825685710238>.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145 – 161, 1985. ISSN 0304-3932. doi: DOI: 10.1016/0304-3932(85)90061-3. URL <http://www.sciencedirect.com/science/article/B6VBW-4D5NXWG-X/2/a312ece0f2adfec0212068ee>

Robert C. Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183, 1973a. ISSN 00058556. URL <http://www.jstor.org/stable/3003143>.

Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):pp. 867–887, 1973b. ISSN 00129682. URL <http://www.jstor.org/stable/1913811>.

Roy Radner. Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica*, 40(2):289–303, 1972. ISSN 00129682. URL <http://www.jstor.org/stable/1909407>.

Richard Roll. A critique of the asset pricing theory's tests part i: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2):129 – 176, 1977. ISSN 0304-405X. doi: DOI: 10.1016/0304-405X(77)90009-5. URL <http://www.sciencedirect.com/science/article/pii/0304405X77900095>.

Bryan R Routledge and Stanley E. Zin. Generalized disappointment aversion and asset prices. *Journal of Finance*, In press, 2011.

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