Liquidity Discovery and Asset Pricing*

Michael Gallmeyer
Carnegie Mellon University

Burton Hollifield
Carnegie Mellon University

Duane Seppi
Carnegie Mellon University

First Draft: April 2003
Current Draft: October 2003

*We thank Robert Engle, Gideon Saar, Stathis Tompaidis, Tan Wang, Robert Whitelaw and seminar participants at the Commodity Futures Trading Commission, Instituto Tecnológico Autónomo De México, New York University, and the University of British Columbia Summer Finance Conference for helpful comments.

Abstract

Asset prices are random, in part, because of uncertainty about the preferences of potential counterparties and their future demands for securities. We call such randomness *liquidity risk*. We model the endogenous dynamics of liquidity risk, the risk premium for bearing liquidity risk, and the role of market trading in the *liquidity discovery* process through which investors form expectations about future liquidity. Our model also provides explanations for "price support levels" and "flights to quality."

1 Introduction

Stocks and other long-dated assets are claims on streams of cash flows that continue long after the likely holding horizons of most investors. Although the value of long-dated securities derives ultimately from their underlying cash flows, their valuations take into account the fact that ownership of these cash flows will change hands over time. Consequently, beliefs about the preferences of future potential trading counter-parties affect investors' current demands for long-dated assets. In particular, these counter-party preferences will lead to the net demands for securities that investors will face when they need to re-trade in the future.

The interaction of prices, order flows, and uncertainty about counter-party preferences is at the intersection of general equilibrium asset pricing and market microstructure. The canonical asset pricing model of Lucas [1978] assumes that investor trading strategies are common knowledge. Prices are random because the arrival of news about future cash flows, but future state-contingent security demands are perfectly anticipated. Since there is no learning about investor preferences, the only learning is about the exogenous underlying cash flows. Moreover, the trading process does not reveal information in a Lucas economy. It simply "digests" information that arrives exogenously from other sources by reallocating securities across investors.

Learning is central in the microstructure approach of Kyle [1985] and Glosten and Milgrom [1985] and in the rational expectations model of Grossman and Stiglitz [1980]. Since informed investors see non-public signals about future cash flows, the trading process itself is informative. However, uninformed investors only learn about cash flows; not about their informed counter-parties' preferences. Indeed, the informed investors' preferences play no direct role in asset pricing once the signal extraction problem is solved in Kyle [1985] and Glosten and Milgrom [1985] or in the uninformed investors' security demands once prices are set in Grossman and Stiglitz (1980).

Our starting point is the idea that future prices are random, in part, because of uncertainty about how other investors' demands for securities will respond to future events and cash flows. As first described in Grossman [1988] and Kraus and Smith [1989], uncertainty about other investors' preferences is in addition to uncertainty about future cash flows. We call uncertainty about counter-party preferences *liquidity risk* and uncertainty about future cash flows *cash flow risk.* The presence of liquidity risk reduces the level of risk sharing in the market, leading to a liquidity risk premium in expected returns.

It is unlikely that personal preferences are common knowledge in full. First, preferences may experience random shocks over time. Indeed, investors may not always perfectly know a priori the future evolution of their own preferences. In this case, neither the investor himself nor his counter-parties know exactly what his demand for securities will be in the future. It is reasonable, however, to suppose that investors learn their personal preferences before their preferences become known to others. Second, an investor may know his own future preferences a priori, but this information may not be common knowledge among his counter-parties. In other words, the investor's preferences appear random to others, but not to himself.

Both intuitions are consistent with our analysis. In either case, investors can learn about their counter-parties' likely demand for securities in the future from the demand these preferences induce for securities today. Thus, trading itself generates random price changes, as in French and Roll [1986], without cash flow shocks. In such an environment, the trading process is both a mechanism for learning about non-public investor preferences as well as for digesting information about cash flows. We call the process of learning about counter-party preferences *liquidity discovery* because investors are uncertain a priori about the prices at which they will find willing counter-parties for trades in the future. Liquidity discovery reduces liquidity risk, but the amount that can be learned from current prices and order flows is itself endogenous.

We develop a model to illustrate both how liquidity risk arises and how trading leads to liquidity discovery. Investors differ in their holding periods. Short-horizon investors trade default-free long-dated bonds with long-horizon investors, but are uncertain about the long-horizon investors' time preferences. In particular, short-horizon investors' cannot always infer the long-term investors' discount rates for individual subperiods. Consequently, the short-horizon investors are uncertain about the price at which they will be able to trade the bond in the future. Our main results are:

- Prices are random because of uncertainty about the long-lived investors' future preferences.
- Liquidity risk is a priced risk factor. Expected bond returns include a liquidity premium in excess of the risk-free rate.
- Preference information is either fully or partially revealed via order flows. Thus, the level of

prevailing liquidity risk is endogenous.

- Market liquidity exhibits endogenous *price support* levels. Small order flow variations can lead to abrupt and even discontinuous jumps (e.g., "crashes") from one price support level to another.

- Trading volume is correlated with the volatility of future liquidity risk and the size of the current risk premium. Sufficient conditions are given for a non-monotone relation between prices, volumes, liquidity risk, and the liquidity risk premium.

- The predicted co-movements of volume, prices, returns and bond/bill spreads are consistent with so-called "flights to quality" such as in the aftermath of the 1998 Russian bond default.

- Unverifiable self-reports of preferences are not incentive compatible and, hence, neither reduce liquidity risk nor Pareto improve investors' welfare.

Liquidity discovery offers a new explanation for the growing body of empirical evidence showing that traditional microstructure variables — such as order flow/price impact coefficients and the probability of informed trading — are "priced" in the sense that they explain expected returns. This relation was documented first in Amihud and Mendelson [1986] and more recently in Easley, Hvidjkaer, and O'Hara [2002] and Stambaugh and Pastor [2002]. O'Hara (2003) appeals to this empirical evidence to argue that the learning process inherent in trading has a first-order impact on investors' portfolio choices and, hence, on required risk premia. However, it is preference uncertainty, rather than asymmetric information about cash flows, that is central in our liquidity discovery explanation.

Over time, investors form expectations about the liquidity available from other investors in the market. In particular, investors care not only about current liquidity, but also about the terms-of-trade they can anticipate in the future. Imperfectly known future liquidity leads to endogenous support and resistance levels in current prices. Surprises in market volume can cause prices to break through prevailing technical price support levels and lead to uncertainty about where the "points of liquidity" will be in the future. The heightened liquidity risk leads, in turn, to higher risk premia until the new support levels are discovered. In this context, "liquidity crises" are simply subsequent realizations of liquidity that are worse than what was expected ex ante.

Our paper builds on previous work by Grossman [1988] and Kraus and Smith [1989]. Grossman [1988] introduced the idea that investors' future trading plans are not common knowledge and that, consequently, prices can change, sometimes dramatically, as the flow of orders into the market reveals the latent strategies investors are following.

Kraus and Smith [1989] is the first formal model of liquidity risk. Their model has multiple "sunspot" equilibria. In one possible equilibrium trade occurs without resolving any uncertainty about counter-party preferences. One contribution of our paper is that we model both the extent to which counter-party information is revealed and the microstructure mechanism through which this happens. We also study the impact of liquidity risk on market risk premia.

Other related work includes Kraus and Smith [1996, 1998] which use counter-party uncertainty to endogenize noise trading in a rational equilibrium model with asymmetric cash flow information. Smith [1993] develops an overlapping generations model in which liquidity risk takes the form of uncertainty about how many traders will arrive in the future. Since traders are risk averse, future prices are random because the aggregate future demand for intertemporal smoothing is random. Leach and Madhavan [1992] and Saar [2001] model dynamic updating of market maker beliefs about investor demand based on observed order flows. Vayanos [1999, 2001] models strategic trading when a large investor needs to trade for non-informational reasons, but can choose the timing of his trades. Liquidity providers seek to learn his preferences — that is, the total number of shares he intends to buy — by observing his trades over time. Pedersen and Acharya [2003] provide an explanation for the Stambaugh and Pastor [2002] results in a model with cross-sectionally correlated transaction costs. In contrast, our approach links the concept of liquidity to the fundamental demand for securities rather than to transactional frictions.

Following the market crash of 1987, a number of papers sought to explain how small variation in order flow might lead to discontinuous changes in price. In Gennotte and Leland [1990] discontinuous price change arises in a rational expectation setting due to confusion about whether traders are informed about future payoffs or about noisy supply shocks. Madrigal [1996] considers the impact of private information about noise trading in a Kyle [1985] model. Madrigal and Scheinkman [1997] has discontinuous price changes in a model of strategic market marker learning with heterogeneously informed agents. In all three papers, small changes in trading volume sometimes lead to large revisions in expectations about future cash flows. In contrast, the abrupt price changes in our model arise due to uncertainty about the endogenous prices at which current counter-parties will be willing to re-trade in the future.

Our paper is organized as follows. Section 2 presents a model of liquidity risk based on uncertainty about the time preferences of a long lived investor. Section 3 is the conclusion. All proofs are in the Appendix.

2 Liquidity risk and time preferences

We consider a pure exchange economy with three trading dates $t = 1, 2, 3$. There is one traded security: A two-period discount bond which pays one unit of consumption at time 3. The bond has no cash flow risk because it is default-free. Although we refer to a single bond, this bond proxies for the entire long-term bond market. Our model is a general equilibrium model of market-wide, or systematic, liquidity risk premia.

Two groups of competitive investors trade with each other. The first is a continuum of identical long-horizon investors, denoted by the subscript L, with three-period preferences

$$u(c_{L1}) + \delta_1 \, u(c_{L2}) + \delta_1 \delta_2 \, u(c_{L3}). \tag{1}$$

The two time-preference parameters δ_1 and δ_2 — for discounting between dates 1 and 2 and between dates 2 and 3 respectively — are known ex ante only to the long-horizon investors.[1] In equilibrium, there is no price randomness from the perspective of the long-horizon investors. They know their time-preferences and there is no cash flow risk. Hence, no expectations are necessary in (1).

The long-horizon investors have individual endowments of $e_{L1} = 0$ of consumption at date 1, $e_{L2} > 0$ units of consumption at date 2, and start out holding all of the long-dated bond. The endowment structure causes the long-horizon investors to sell some of their bond holdings at date 1. In particular, e_{L2} cannot be traded directly at date 1 either because of moral hazard or because ownership is unverifiable ex ante.

The second group is a continuum of identical short-horizon investors, denoted by the subscript

[1] We ensure that the long-horizon investors follow time-consistent consumption plans by solving their optimization problem through backward recursion. For a discussion of time-consistency with such preferences, see Zin [2001].

S, who have expected utility preferences over consumption at dates 1 and 2:

$$v(c_{S1}) + \beta\, \mathrm{E}_{S1}[v(c_{S2})], \tag{2}$$

where E_{S1} denotes the short-horizon investors' expectations given the information available to them at date 1 and v is increasing, concave, differentiable and satisfies the Inada conditions. The short-horizon investor's preferences are common knowledge and they have individual endowments of $e_{S1} > 0$ units of consumption at time 1 and $e_{S2} \geq 0$ unit of consumption at date 2.

The only uncertainty is that short-horizon investors do not know the long-horizon investors' time-preferences δ_1 and δ_2 for the two subperiods. The short-horizon investors' priors over δ_1 and δ_2 are given by the distribution:

$$\mathrm{Prob}\,(\delta_1 \leq x, \delta_2 \leq y) \equiv F(x, y). \tag{3}$$

We assume δ_1 and/or δ_2 (i.e., at least one) has a continuous distribution to keep the equilibrium from being generically fully revealing as in Radner [1972].

The short-horizon investors use prices and trades at date 1 to learn about the long-horizon investors' time-preferences. We call this *liquidity discovery.* In particular, the bond price at date 2 is uncertain because it depends on δ_2. If trading at date 1 fully reveals δ_2, then the bond is riskless for the short-horizon investors (as well as for the long-horizon investors) between dates 1 and 2. If δ_2 cannot be inferred at date 1, then short-horizon investors are uncertain about the bond price at date 2. In this case, holding the bond exposes them to *liquidity risk.*

Long-horizon investors never experience any liquidity risk. From their perspective, the long-horizon bond is riskless between dates 1 and 2 as well as between dates 2 and 3. In a sense, they are like the informed investors in Grossman and Stiglitz (1980) but with the significant difference that the component of future prices about which they are informed arises — not from private information about an exogenous cash flow — but rather endogenously from their own future actions (as induced by their preferences). Hence, trading at date 1 affects not only how much information the "uninformed" short-horizon traders learn through liquidity discovery, but also the future price itself via the impact of the long-horizon investor's date 1 portfolio holdings on their date 2 trades.

The motive for trade in this economy is intertemporal consumption smoothing. However, if the short-horizon investors cannot perfectly anticipate the long-horizon investors' future security demands, then the resulting liquidity risk prevents investors from smoothing their consumption perfectly.

A time line for the economy is in Figure 1. Let P_1 and P_2 be the prices of the bond at times 1 and 2. Long-horizon investors trade the bond to shift consumption from date 2 to date 1. In particular, given their zero consumption endowment e_{L1}, they sell some of their bonds at P_1 to finance consumption c_{L1} at date 1. Let θ_{L1} denote the number of bonds the long-horizon investors hold at date 1 so that $1 - \theta_{L1}$ is the number of bonds they sell per capita to the short-horizon investors. Later at date 2 they use their endowment e_{L2} to buy back $\theta_{L2} - \theta_{L1}$ bonds at price P_2 to bring their total bond holdings up to θ_{L2}. In equilibrium, the long-horizon investors buy back *all* of the bonds — implying $\theta_{L2} = 1$ and $\theta_{L2} - \theta_{L1} = 1 - \theta_{L1}$ — since the short-horizon investors do not value consumption at time 3. Conversely, the short-horizon investors increase their consumption at time 2 by buying bonds at time 1 and then reselling them at time 2. The more general point here is that implementing optimal consumption sharing rules in incomplete markets usually requires investors to follow dynamic trading strategies in which they switch their roles as buyers and sellers over time.

We show below that, in equilibrium, the resale price P_2 depends on the net trade $1-\theta_{L1}$ at date 2 and on the long-horizon investors' second-period time-preference δ_2. The short-horizon investors may not always learn the long-horizon investors' time preference parameters from the date 1 price and allocation. The bond price P_2 may, therefore, be a random variable from the short-horizon investors' point of view. Let π be the short-horizon investors' beliefs at date 1 about the distribution of P_2.

2.1 Equilibrium

A rational expectations equilibrium consists of:

- Share allocations $\theta_{L1}(\delta_1, \delta_2)$ and $\theta_{L2}(\delta_1, \delta_2)$ as functions of the long-horizon investors' time-preference parameters.

- Bond prices $P_1(\delta_1, \delta_2)$ and $P_2(\delta_1, \delta_2)$ as functions of the long-horizon investors' time preference parameters:

- Posterior beliefs $\pi(P_2|\theta_{L1}, P_1)$ for the short-horizon investors about the distribution of the date 2 bond price conditional on the bond allocation and price the date 1.

that satisfy:

- Optimality: Both the long-horizon and the short-horizon investors' portfolios are optimal given the prices.
- Rational expectations: The short-horizon investors beliefs about P_2 satisfy rational expectations.
- Market clearing: Supply equals demand for the bonds.

We now turn to characterizing the equilibrium. Using the budget constraints

$$\begin{aligned} c_{L1} &= P_1(1-\theta_{L1}), \\ c_{L2} &= e_{L2} - P_2(\theta_{L2} - \theta_{L1}), \\ c_{L3} &= \theta_{L2}, \end{aligned} \tag{4}$$

the portfolio problem for the long-horizon investor is

$$\max_{\theta_{L1},\theta_{L2}} u(P_1(1-\theta_{L1})) + \delta_1\, u(e_{L2} - P_2(\theta_{L2} - \theta_{L1})) + \delta_1\, \delta_2\, u(\theta_{L2}). \tag{5}$$

The first-order conditions for their optimal bond holdings are:

$$P_1 = \frac{\delta_1\, u_c(e_{L2} - P_2(\theta_{L2} - \theta_{L1}))P_2}{u_c(P_1(1-\theta_{L1}))} \tag{6}$$

$$P_2 = \frac{\delta_2 u_c(\theta_{L2})}{u_c(e_{L2} - P_2(\theta_{L2} - \theta_{L1}))}. \tag{7}$$

The portfolio problem for the short-horizon investors is

$$\max_{\theta_{S1}} v(e_{S1} - P_1\theta_{S1}) + \beta\, E_{S1}\left[v\left(\theta_{S1}\tilde{P}_2 + e_{S2}\right)\right], \tag{8}$$

where expectations over $\tilde{P}_2$ are taken using the short-horizon investors' beliefs. Since they do not value consumption at date 3, they sell of all their bond holdings at date 2 to finance consumption. The first-order condition for their optimal bond position is:

$$P_1 v_c(e_{S1} - P_1\theta_{S1}) = \beta\, E_{S1}\left[v_c(\theta_{S1}\tilde{P}_2 + e_{S2})\tilde{P}_2\right]. \tag{9}$$

Market clearing requires that the long-horizon investors buy back all of the bonds from the short-horizon investors at date 2. Imposing market-clearing, $\theta_{L2} = 1$, and substituting (7) into (6), gives:

$$\begin{aligned} P_1 &= \frac{\delta_1\delta_2\, u_c(1)}{u_c(P_1(1-\theta_{L1}))} && (10)\\ P_2 &= \frac{\delta_2 u_c(1)}{u_c(e_{L2} - P_2(1-\theta_{L1}))}. && (11) \end{aligned}$$

Note that P_2 only depends on the net trade $1-\theta_{L1}$ at date 2 and on the long-horizon investor's time-preference δ_2 between dates 2 and 3. Since the date 2 net trade is perfectly predictable — it exactly offsets the net trade at date 1 — the only reason P_2 is random is because of uncertainty about δ_2. In particular, P_2 is higher if the long-horizon investors are more patient (i.e., if δ_2 is large) and lower if they are impatient (i.e., if δ_2 is smaller). Thus, (11) confirms that holding the bond is riskless for the long-horizon investors but potentially risky for the short-horizon investors unless they can infer δ_2 from market conditions at date 1.

Liquidity discovery is the process through which the short-horizon investors learn about the terms of trade they will face at date 2 when selling back the bond to the long-horizon investors. In particular, substituting the market-clearing bond price P_1 and volume $1-\theta_{L1}$ into the date 1 first-order condition for the long-horizon investors (10) lets the short-horizon investors compute a summary statistic for the long-horizon investor's *cumulative* time-preference between dates 1 and 3

$$z = \frac{P_1 u_c(P_1(1-\theta_{L1}))}{u_c(1)} = \delta_1\delta_2. \tag{12}$$

The subperiod time-preference δ_2 is fully revealed if only one of the possible δ_2's could have led to the observed z (i.e., given $F(\delta_1, \delta_2)$). If not, then δ_2 is not fully revealed by the market-clearing price and trading volume at date 1.

Lemma 1 *The short-horizon investors' equilibrium beliefs about the long-horizon investors' second-period time-preferences are:*

$$Prob\,(\,\delta_2 \leq x |\, \theta_{L1}, P_1) = Prob\,(\,\delta_2 \leq x |\, \delta_1 \delta_2 = z)\,, \tag{13}$$

The short-horizon investors' beliefs about P_2 follow from Lemma 1 and the second-period market clearing price function in equation (11).

A useful distinction here is the difference between *local preferences* and *global preferences.* Market data only reveal information about those aspects of investor preferences that are relevant for the marginal demand and supply of securities under current market conditions. In general, however, prices and volumes at a single date do not always reveal enough information to infer investor preferences under *all* possible future conditions. For example, the statistic $z = \delta_1 \delta_2$ summarizes the long-horizon investors' local preferences as as they matter for prices and trades at date 1, but it is not always possible for δ_1 and δ_2, the parameters that determine the long-horizon investors global preferences, to be identified separately.

Lemma 2 *The set of types (δ_1, δ_2) that pool in equilibrium is independent of the other parameters of the economy.*

Existence of a rational expectations equilibrium is established by showing that, for any possible realization of the long-horizon investors' time-preference parameters, the market-clearing conditions can be solved.

Proposition 1 *Under the previously stated assumptions concerning preferences and endowments, an equilibrium exists. If the long-horizon investors have an elasticity of intertemporal substitution greater than or equal to one, the equilibrium is also unique.*

The restriction on the elasticity of intertemporal substitution simply ensures that consumption at date 2 is a normal good and, hence, that the long-horizon investor's supply curve slopes up — that is, rather than sloping down and, thus, potentially crossing the downward sloping short-horizon investor's demand curve more than once.

2.2 Examples

We next present several analytic and numerical examples to illustrate the workings of the model more concretely. Some of the intuitions from these examples are formalized in Section 2.3.

Analytic results with logarithmic preferences. When investors have logarithmic preferences $u(c) \equiv v(c) \equiv \ln(c)$, we can solve for the equilibrium in closed-form. The log case is special in that the long-horizon investors' net trade at date 1 does not depend on the bond price P_1. From (10) and (11), the date 1 net trade is

$$1 - \theta_{L1} = \frac{1}{\delta_1 \delta_2} \tag{14}$$

and the date 2 market-clearing bond price is

$$P_2 = \frac{\delta_1 \delta_2 \, e_{L2}}{1 + \delta_1}. \tag{15}$$

The corresponding consumption profile of the long-horizon investor is

$$c_{L1} = \frac{P_1}{\delta_1 \delta_2}, \tag{16}$$

$$c_{L2} = e_{L2} \left(\frac{\delta_1}{1 + \delta_1} \right). \tag{17}$$

Given (14), the short-horizon investors can infer $\delta_1 \delta_2$ from the date 1 trading volume alone without reference to the price P_1. Letting $E[\cdot|\delta_1 \delta_2]$ denote the short-term investor's expectations conditional on $\delta_1 \delta_2$ (from Lemma 1), we can write the equilibrium bond price at date 1 as

$$P_1 = \frac{e_{S1} e_{L2} \, \delta_1 \delta_2 \, E\left[\left(e_{S2}(1 + \delta_1) + e_{L2}\right)^{-1} \Big| \, \delta_1 \delta_2 \right]}{\frac{1+\beta}{\beta} - e_{S2} \, E\left[\left(e_{S2} + \frac{e_{L2}}{1+\delta_1} \right)^{-1} \Big| \, \delta_1 \delta_2 \right]}. \tag{18}$$

In order to compute a risk premium for liquidity risk, we also calculate the short-horizon investor's shadow price for a one-period bill paying one unit of risk-free consumption at date 2

$$b_1 = \frac{e_{S1} \, E\left[(e_{S2} + \frac{e_{L2}}{1+\delta_1})^{-1} \Big| \, \delta_1 \delta_2 \right]}{\frac{1+\beta}{\beta} - e_{S2} \, E\left[\left(e_{S2} + \frac{e_{L2}}{1+\delta_1} \right)^{-1} \Big| \, \delta_1 \delta_2 \right]}. \tag{19}$$

Including an actual tradable one-period bill along with the bond would eliminate all liquidity risk in this simple setting. Given two securities and only one source of preference uncertainty, the short-lived investors could generically recover the long-horizon investor's full preferences. Liquidity risk with multiple securities requires multiple dimensions of preference uncertainty.[2] The advantage of our current model is that it leads to a particularly simple illustration of liquidity risk.

Numerical results for log preferences and uniform priors. We specialize the log model by assuming that δ_2 is binomially distributed with a state space $\{\delta_I, \delta_P\}$ where $\delta_I < \delta_P$. If δ_2 is δ_I, we say the long-lived investors are *impatient* at date 2. If δ_2 is δ_P, they are *patient.* Given $\delta_2 = \delta_I$, the date 1 time-preference δ_1 has a continuous conditional density f_I defined on an interval support. Given $\delta_2 = \delta_P$, the conditional density f_p for δ_1 is also continuous on an interval support.

We assume patient and impatient investors are equally likely so that $\text{Prob}(\delta_2 = \delta_I) = \text{Prob}(\delta_2 = \delta_P) = 0.5$. For the patient investors, $\delta_P = 1.2$ and δ_1 is uniformly distributed on the interval [0.74, 1.1]. For the impatient investors, $\delta_I = 0.85$ and δ_1 is uniformly distributed on [1.05, 1.3]. The short-horizon investors' time-preference is $\beta = 1.2$. "Discount" factors larger than 1 just mean that later consumption has a high value relative to early consumption. At time 1, the short-horizon investors have an initial endowment of $e_{S1} = 1.4$. At time 2, the long-horizon investors' unsecuritized endowment is $e_{L2} = 1.0$ and the short-horizon investors' endowment is $e_{S2} = 0.6$.

Given our assumptions, the support for $z = \delta_1\delta_2$ conditional on δ_I is nested inside the support for z conditional on δ_P. Thus, there is a critical value z' such that when $z < z'$ the equilibrium outcome is only partially revealing. In other words, there are δ_1's such that both the patient or impatient types can have cumulative time preferences $\delta_1\delta_2 < z'$. However, above z' the equilibrium outcome is fully revealing. Given our numerical assumptions, only δ_P types can have cumulative time-preferences $\delta_1\delta_2 > z'$. In this example the probability of a non-fully revealing equilibrium outcome is 60 percent.

With log preferences, it is more intuitive to work with trading volume $1 - \theta_{L1} = 1/z$ rather than directly with z. Indeed, the model has a natural market microstructure interpretation. The inelastic bond supply from the long-horizon investors at time 1 can be seen as a market order.

[2] In addition, cash flow risk or some other type of randomness is needed to avoid arbitrage between the riskless bills and the otherwise riskless (for the better-informed long-horizon investors) bond.

Similarly, the short-horizon traders' first-order condition at date 1 characterizes a market liquidity supply function $P_1(\cdot)$ which gives the price P_1 at which they are willing to absorb $1/z$ bonds.

Figures 2a and 2b are an example of a (δ_1, δ_2) pair where the equilibrium is fully revealing. Figure 2a shows supply and demand at time 1. The solid downward sloping curve is the short-horizon investors' demand when they know they are trading with a patient δ_P long-horizon investor. The dash-dot-dot (- · ·) curve is the short-horizon investor's demand if they know they are trading with an impatient δ_I investor. The dashed (- - -) curve in between is their demand when they are uncertain about whether they are trading with patient or impatient investors. The solid vertical line is the supply curve of a long-horizon investor who happens, in this figure, to be both patient (δ_P) and has a high initial time-preference δ_1. Her supply curve is inelastic due to logarithmic utility. The dashed vertical line at $1/z'$ is the minimal trade for an impatient δ_I investor given her highest possible date 1 time-preference δ_1 (i.e., given the density f_I). Since short-horizon investors know that impatient investors never sell fewer than $1/z'$ bonds, the realized trade $1/z < 1/z'$ fully reveals the long-horizon investors as being the δ_P type.

Figure 2b illustrates supply and demand at time 2. Since the short-horizon investor inelastically sells all of her bonds, it is now her supply curve that is vertical. The two downward sloping curves are the long-horizon investors' demands for bonds conditional on being patient or impatient. In this case the upper curve is the relevant schedule because the long-horizon investor is patient.

Figures 3a and 3b show a different (δ_1, δ_2) pair for which the equilibrium outcome is only partially revealing. Figure 3a again shows supply and demand at time 1. The realized supply (i.e. the solid vertical line) could come from either (a) a patient investor with a *low* initial time-preference δ_1' or from (b) an impatient investor with a sufficiently *large* initial time-preference δ_1'' such that $\delta_1'\delta_P = \delta_1''\delta_I$. Consequently, the short-horizon investor cannot tell whether she is trading against a patient or an impatient investor. The dashed (- - -) line is her demand curve given her uncertainty about the long-horizon investor's type. In Figure 3b, the short-horizon investor inelastically sells all of her bonds from time 1. However, since the date 1 volume did not fully reveal the long-horizon investor's type, two prices are possible at date 2 corresponding to the two possible investor types.

Figure 4a plots the bond price P_1 and the expected price $\mathrm{E}_{S1}[P_2]$ versus the trading volume $1/z$ at date 1. A small net trade $1/z < 1/z'$ fully reveals that the long-horizon investors are the patient δ_P type. In this case, short-horizon investors know (from (15)) that P_2 will be $\frac{z\, e_{L2}}{1+z/\delta_P}$ with certainty and, hence, are willing to pay a high price P_1 for the bond at time 1. Prices are decreasing in trading volume because more selling at date 1 means that the long-horizon investors must buy back more bonds at date 2 which depresses P_2. Since the net trade $1 - \theta_{L1}$ is publicly observable at time 1, the "market overhang" effect at date 2 is fully anticipated when P_1 is determined at date 1. At the critical volume $1/z'$, both P_1 and $\mathrm{E}_{S1}[P_2]$ have discrete jumps downward. This is because volumes $1/z > 1/z'$ are consistent with either δ_P and δ_I. Now there is liquidity risk.

Turning to the pricing of liquidity risk, Figure 4b shows that the standard deviation of the bond return is roughly 10 percent in non-fully revealing states. Thus, our example is consistent with the positive volume/volatility relation seen empirically (see Karpoff [1987]). Figure 4c contrasts the shadow risk-free return on one-period bills with the expected return on two-period bonds. The two are equal when the equilibrium outcome is fully revealing, but there is a small but non-trivial liquidity premium in the bond's expected return when the equilibrium outcome is only partially revealing. The liquidity premium is roughly 60 basis points from Figure 4d.

The large swings in the shadow one-period interest rate in Figure 4c follow from the changing allocations of consumption associated with different levels of trading at date 1. Low volumes $1/z$ fully reveal that the long-horizon investors are patient. As long-horizon investor sell more bonds, short-horizon investors forgo more date 1 consumption and have more date 2 consumption. This raises both the one-period interest rate and the return on the long-dated bond. At the critical volume $1/z'$, the date 1 consumption of the short-horizon investors jumps up — since they now pay less for essentially the same number of bonds due to the possibility that the long-horizon investor is impatient — and their expected date 2 consumption falls — since impatient investors will, on average, pay less to buy back the bonds — thereby causing both the one-period spot rate and the expected bond return to fall. Thus, large volumes at date 1 increase the likelihood that the long-horizon investors are impatient and, hence, are unfavorable for the short-horizon investors' welfare.

These patterns of prices, returns, and volumes offer a natural explanation for the "flight to quality" that followed the 1998 Russian default. Consider the market for non-Russian long-dated debt. Short-horizon investors — investors currently not holding long-dated bonds — know that long-horizon investors with large bond positions (e.g., pensions, banks, hedge funds) will react to major bond defaults by selling off bonds. This sell-off is due to some combination of short-term liquidity needs and longer-term views about bond market fundamentals. In particular, the default might cause long-horizon investors to become fundamentally more bearish about long-dated debt — that is, their future (i.e., date 2) time preferences will be impatient — leading to low future demand for bonds. Alternatively, they might stay fundamentally bullish — or, in our terminology, they are patient — and their future bond demand will bounce back once the current crisis passes.

If the current bond sell-off is small, this fully reveals that the long-horizon investors are fundamentally bullish and that the sell-off is due to transitory technical conditions in the market. However if the sell-off is large, then short-horizon investors cannot tell whether the long-horizon investors are bullish with large short-term liquidity needs or bearish with small short-term liquidity needs. Consequently, long-dated bond prices fall (due to their lower anticipated future resell prices) and short-term bill prices rise (due to the shift between current and future expected consumption). Spreads between short-term treasuries and long-dated credit sensitive debt rises due to the liquidity risk at time 2 given the uncertainty about the future preferences of long-horizon investors.[3]

The interplay between trading and liquidity discovery is particularly stark in the log model due both to the price-insensitivity of the long-horizon investor's net demand at date 1 and the uniform/binomial distribution over the subperiod time-preferences (δ_1, δ_2). Our next two examples relax these restrictive features.

Example with Beta priors. In this example we keep the log-log preferences but now assume that the short-term investors' priors over δ_1 have an extended beta (2,2) distribution scaled to cover the same conditional supports for δ_1 given δ_P and δ_I. Otherwise, the parameter values are unchanged from the first example.

Figure 5 has the same layout as Figure 4. The main point is that the market at date 1 still has an endogenous price support level in Figure 5a. The only difference is that, with beta priors, the transition between the fully revealing region and the partially revealing region is continuous; unlike the discontinuous jump with the uniform priors in Figure 4a.

[3]The story is told here in terms of time-preferences to be consistent with the current model, but a similar story could be told with cash flow risk and uncertainty about investor's willingness to bear credit risk.

Example with square-root preferences and uniform priors. In this example, the long-horizon and short-horizon investors have constant relative risk aversion utility $u(c) = c^{1-\gamma}/(1-\gamma)$ with a coefficient of relative risk aversion $\gamma = 0.5$. The other parameters are unchanged from the first uniform example except that the short-horizon investor's first-period consumption endowment is $e_{S1} = 1.35$.

With square-root preferences, both the volume $1 - \theta_{L1}$ and price P_1 are needed for the liquidity discovery process. We see this in Figure 6a where there are multiple prices for some volumes. This is because, from the long-horizon investor's first-order condition (10), the summary statistic $z = \delta_1\delta_2 = \sqrt{P_1/(1-\theta_{L1})}$. Thus, unlike the inelastic log bond supply, here the supply of bonds at date 1 is increasing in P_1. Now, rather than a minimal volume $1/z'$, there is a minimal (i.e., "left most") positively-sloped bond supply schedule for impatient investors. Hence, long-horizon investors who trade price/volume pairs $(P_1, 1-\theta_{L1})$ on the upper part of the short-horizon investors demand schedule to the left of the minimal impatient bond supply curve are fully revealed as the patient type. In Figure 6a the right-most point on the upper demand curve at date 1 (used when trading with fully revealed patient investors) and the left most point on the lower demand curve (when the long-horizon investor's type is not fully revealed) are connected by this minimal bond supply schedule.

The square-root specification provides a clear illustration of the distinction between global and local preferences. Consider the *Marshallian* net demand for the long-horizon investor at time 1 holding P_2 fixed. This is in contrast to the demand from the composite first-order condition (10) which implicitly allows P_2 to vary so as to preserve (7). Figure 7 plots one pair of Marshallian net demands $g(P_1; P_2 = P_2(\delta_P, \theta_{L1}))$ for the patient investors holding P_2 equal to its equilibrium value $P_2(\delta_P, \theta_{L1})$ and $g(P_1; P_2 = P_2(\delta_I, \theta_{L1}))$ for the impatient investor holding P_2 equal to $P_2(\delta_I, \theta_{L1})$. With partial revelation, the two Marshallian demand curves intersect at the equilibrium price/volume pair. Thus, in the rational expectations equilibrium the short-horizon investor knows that, locally, the long-horizon investors preferences are such that his demand curve passes through the market-clearing price/volume pair. However, she cannot tell what the long-horizon investor's Marshallian demand looks like globally away from this point.

Example with power short-horizon utility. In Figure 8 the long-horizon investors again have square-root preferences but now the short-horizon investors have constant relative risk aversion equal to 2. The distributional assumptions and endowments are the same as in the first log/uniform case. The example is a numerical illustration of the comparative static that increasing the short-horizon investor's risk aversion increases the required risk premium for liquidity risk.

2.3 General Properties

The way the supports for the cumulative time-preference $\delta_1\delta_2$ overlap and are nested conditional on different second-period time-preferences δ_2 is critical to the co-movement of prices, returns, volatility, and risk-premia with the date 1 volume. The roughly monotone patterns (i.e., up to the effect of volume) in our numerical examples followed directly from the particular nesting assumed there. However, it is easy to construct distributions $F(\delta_1, \delta_2)$ for which the predicted relationships would non-monotone.

Proposition 2 *If, given the distribution $F(\delta_1, \delta_2)$, the sets of possible products $\delta_1\delta_2$ associated with multiple values of δ_2 are surrounded by (or, alternatively, surround) the sets of possible products $\delta_1\delta_2$ which are uniquely associated with single values of δ_2, then the resulting prices P_1, expected prices $\mathrm{E}_{S1}[P_2]$, expected bond returns, return volatilities and liquidity risk premia are non-monotone in z.*

The existence of price support levels also generalizes. As the beta example illustrates:

Proposition 3 *Abrupt changes in the conditional density $f(\delta_1|\delta_2)$ for some value of δ_2 relative to the conditional densities for other pooled δ_2s lead to abrupt changes in prices, spreads, volatility, and risk premia relative to z.*

Liquidity risk prevents investors from smoothing their consumption perfectly. Thus, it is natural to wonder whether the long-horizon investors can improve their welfare by simply announcing their type before the first round of trade.

Proposition 4 *If the long-horizon investor is patient, then her utility would be higher if her type could be credibly pre-announced. If the long-horizon investor is impatient, then her utility would be lower if her type were pre-announced.*

Unverifiable statements by long-horizon investors about their preferences are not credible. Impatient investors benefit from being confused with patient investors at date 1 because date 1 bond prices are higher — due to their higher anticipated date 2 resell value with patient investors — than if it were common knowledge that they are impatient. As a consequence, long-horizon investors cannot remove liquidity risk by simply announcing their type, unless the announcements are verifiable.

3 Conclusion

This paper provides a bridge between general equilibrium asset pricing and market microstructure as recently called for in O'Hara (2003). Our model is based on ideas of liquidity risk and liquidity discovery. In some states investors can fully infer the future demands of their other investors from current market data while in others they cannot. When the preferences of potential counterparties to future trades are uncertain, the resulting randomness in the future resale prices of long-dated securities exposes investors to liquidity risk. Accordingly, a risk premium is required to clear the market for the long-dated securities with liquidity risk. Thus, our model provides a possible explanation for why market microstructure variables (such as order flows and price impacts) empirically predict asset pricing variables such as future price volatility and risk premia.

The only stochastic shocks in our model are to the preferences of the long-term traders — the initial liquidity demanders. The only risk that liquidity suppliers face is liquidity risk. In current work we are extending our model to allow for both cash flow risk and liquidity risk. How much of an additional risk premium would liquidity suppliers require to compensate them for liquidity risk in the presence of cash flow risk? How large is the liquidity risk premium relative to the cash flow risk premium?

References

Amihud, Y. and H. Mendelson, 1986, Asset Pricing and the Bid–Ask Spread, *Journal of Financial Economics*, 17, 223–249.

Easley, D., S. Hvidjkaer, and M. O'Hara, 2002, Is Information Risk a Determinant of Asset Prices?, *Journal of Finance*, forthcoming.

French, K., and R. Roll, 1986, Stock Return Variances: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics.*

Gennotte, G., and H. Leland, 1990, Market Liquidity, Hedging, and Crashes, *American Economic Review*, 999-1021.

Glosten, L., and P. Milgrom, 1985, Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders, *Journal of Financial Economics*, 21, 123–144.

Grossman, S., 1988, An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies, *Journal of Business*, 61, 275–298.

Grossman, S., and J. Stiglitz, 1980, On the Impossibility of Informationally Perfect Markets, *American Economic Review*,70, 393-408.

Karpoff, J., 1987, The Relation between Price Changes and Trading Volume: A Survey, *Journal of Financial and Quantitative Analysis*, 22, 109–126.

Kraus and Smith, 1989, Market Created Risk, *Journal of Finance*, 44, 557-569.

Kraus, A., and M. Smith, 1996, Heterogeneous Beliefs and the Effect of Replicable Options on Asset Prices, *Review of Financial Studies*, 9, 723-756.

Kraus, A., and M. Smith, 1998, Endogenous Sunspots, Pseudo-Bubbles, and Belief and Beliefs, *Journal of Financial Markets,* 1, 151-174.

Kyle, A., 1985, Continuous Auctions and Insider Trading, *Econometrica*, 53, 1315–1335.

Leach, J. and A. Madhavan, 1992, Intertemporal Price Discovery by Market Makers: Active versus Passive Learning, *Journal of Financial Intermediation*, 2, 207-235.

Lucas, R., 1978, Asset Prices in an Exchange Economy, *Econometrica*, 46, 1429–1446.

Madrigal, V., 1996, Non-Fundamental Speculation, *Journal of Finance*, 51, 553-578.

Madrigal, V. and J. Scheinkman, 1997, Price Crashes, Information Aggregation, and Market Making, *Journal of Economic Theory*, 75, 16-63.

Mas-Colell, A., M. D. Whinston, and J. R. Green, 1995, *Microeconomic Theory.*

O'Hara, M., 2003, Liquidity and Price Discovery, *Journal of Finance* 58, 1335-1354.

Pastor, L., and R. Stambaugh, 2002, Liquidity Risk and Expected Stock Returns, *Journal of Political Economy*, forthcoming.

Pedersen, L., and V. Acharya, 2003, Asset Pricing with Liquidity Risk, working paper, NYU.

Radner, R., 1972, Existence of Plans, Prices, and Price Expectations in a Sequence of Markets, *Econometrica*, 44, 289-303.

Saar, G., 2001, Investor Uncertainty and Order Flow Information, working paper, NYU.

Smith, R., 1993, Market Risk and Asset Prices, *Journal of Economic Dynamics and Control*, 17, 555-569.

Vayanos, D., 1999, Strategic Trading and Welfare in a Dynamic Market, *Review of Economic Studies*, 66, 219-245

Vayanos, D., 2001, Strategic Trading in a Dynamic Noisy Market, *Journal of Finance*, 56, 131-171.

Zin, S., 2002, Are behavioral asset-pricing models structural?, *Journal of Monetary Economics*, 49, 215-228.

Appendix

Proof of Lemmas 1 and 2. The results follows directly from equation (10). ■

Proof of Proposition 1. Under the assumptions placed on preferences and endowments, an equilibrium exists by applying a standard fixed point argument to the aggregate excess demand function. See for example Chapter 17 of Mas-Colell, Whinston, and Green [1995].

To show uniqueness of an equilibrium, assume the elasticity of intertemporal substitution of the long-horizon investors is greater than or equal to one. The equilibrium is unique if the time 1 demand and supply curves only cross at one point. A sufficient condition is to show that the supply of the long-horizon investors is increasing in P_1 and the demand curve of the short-horizon investors is strictly decreasing in P_1.

First, examine the demand curve of the long-horizon investors at date 2. From (11), the sign of $\frac{\partial P_2}{\partial(1-\theta_{L1})}$ can be computed by the implicit function theorem:

$$\frac{\partial P_2}{\partial(1-\theta_{L1})} = \frac{P_2^2 u_{cc}(e_{L2} - P_2(1-\theta_{L1}))}{u_c(e_{L2} - P_2(1-\theta_{L1})) - P_2(1-\theta_{L1})u_{cc}(e_{L2} - P_2(1-\theta_{L1}))} < 0$$

given the long-horizon investor preferences are strictly risk averse. In particular, the equilibrium price P_2 is uniquely determined for a particular net trade $1-\theta_{L1}$.

From (10), the slope of the long-investors' time 1 supply curve can be determined again by an application of the implicit function theorem:

$$\begin{aligned}
\frac{\partial(1-\theta_{L1})}{\partial P_1} &= -\frac{u_c(P_1(1-\theta_{L1})) + P_1\,(1-\theta_{L1})\,u_{cc}(P_1(1-\theta_{L1}))}{P_1^2\,u_{cc}(P_1(1-\theta_{L1}))} \\
&= -\frac{u_c(P_1(1-\theta_{L1}))\left(1 + P_1\,(1-\theta_{L1})\,\frac{u_{cc}(P_1(1-\theta_{L1}))}{u_c(P_1(1-\theta_{L1}))}\right)}{P_1^2\,u_{cc}(P_1(1-\theta_{L1}))} \\
&= -\frac{u_c(P_1(1-\theta_{L1}))}{P_1^2\,u_{cc}(P_1(1-\theta_{L1}))}\left(1 + P_1\,(1-\theta_{L1})\,\frac{u_{cc}(P_1(1-\theta_{L1}))}{u_c(P_1(1-\theta_{L1}))}\right).
\end{aligned}$$

Given the first term above is positive, the sign of $\frac{\partial(1-\theta_{L1})}{\partial P_1}$ is determined by the sign of the second term. Since the elasticity of intertemporal substitution is greater than or equal to one, this implies the last term is greater than or equal to zero implying the supply curve at date 1 is increasing or $\frac{\partial(1-\theta_{L1})}{\partial P_1} \geq 0$.

Turning to the date 1 demand function of the short-horizon investors and imposing market clearing $\theta_{S1} = 1 - \theta_{L1}$, its slope is computed as

$$\frac{\partial(1-\theta_{L1})}{\partial P_1} = \frac{P_1(1-\theta_{L1})v_{cc}(\cdot) - v_c(\cdot)}{-P_1^2 v_{cc}(\cdot) - \beta E_{S1}\left[\frac{\partial \tilde{P}_2}{\partial(1-\theta_{L1})}v_c(\cdot) + \tilde{P}_2 v_{cc}(\cdot)\left[\frac{\partial \tilde{P}_2}{\partial(1-\theta_{L1})}(1-\theta_{L1}) + \tilde{P}_2\right]\right]}$$

where the consumptions at each date have been suppressed for clarity. Given the numerator in the above equation is strictly negative and $\frac{\partial \tilde{P}_2}{\partial(1-\theta_{L1})} < 0$, the demand curve is downward sloping $\left(\frac{\partial(1-\theta_{L1})}{\partial P_1} < 0\right)$ if

$$\frac{\partial \tilde{P}_2}{\partial(1-\theta_{L1})}(1-\theta_{L1}) + \tilde{P}_2 \geq 0.$$

Given $\frac{\partial P_2}{\partial(1-\theta_{L1})} = \frac{P_2^2 u_{cc}(e_{L2}-P_2(1-\theta_{L1}))}{u_c(e_{L2}-P_2(1-\theta_{L1}))-P_2(1-\theta_{L1})u_{cc}(e_{L2}-P_2(1-\theta_{L1}))}$,

$$\begin{aligned}\frac{\partial \tilde{P}_2}{\partial(1-\theta_{L1})}(1-\theta_{L1}) + \tilde{P}_2 &= \frac{\tilde{P}_2^2 u_{cc}(\cdot)(1-\theta_{L1})}{u_c(\cdot) - \tilde{P}_2(1-\theta_{L1})u_{cc}(\cdot)} + \tilde{P}_2 \\ &= \frac{\tilde{P}_2 u_c(\cdot)}{u_c(\cdot) - \tilde{P}_2(1-\theta_{L1})u_{cc}(\cdot)} > 0,\end{aligned}$$

establishing that the equilibrium is unique. ■

Proof of Proposition 2. The result follows from the fact that sets of $\delta_1\delta_2$ products that pool different δ_2 have liquidity risk whereas sets of $\delta_1\delta_2$ products that are uniquely associated with single values of δ_2 have no liquidity risk. ■

Proof of Proposition 3. The result follows from Bayes' Rule. ■

Proof of Proposition 4. The long-term investor's consumption profile in either the fully revealing or partially revealing equilibrium is

$$\begin{aligned} c_{L1} &= \frac{P_1}{\delta_1\,\delta_2}, \\ c_L(\delta) &= e_{L,2}\left(\frac{\delta_1}{1+\delta_1}\right). \end{aligned}$$

The investor's date 2 consumption is invariant to whether her type is revealed at date 1. However, her date 1 consumption is increasing in the date 1 price of the bond P_1.

To prove the proposition, it is sufficient to show that a patient long-horizon investor faces a lower P_1 in a partially revealing equilibrium and an impatient investor faces a higher P_1 in a partially revealing equilibrium. This result is immediate from the first order condition of the short-horizon investors. ■

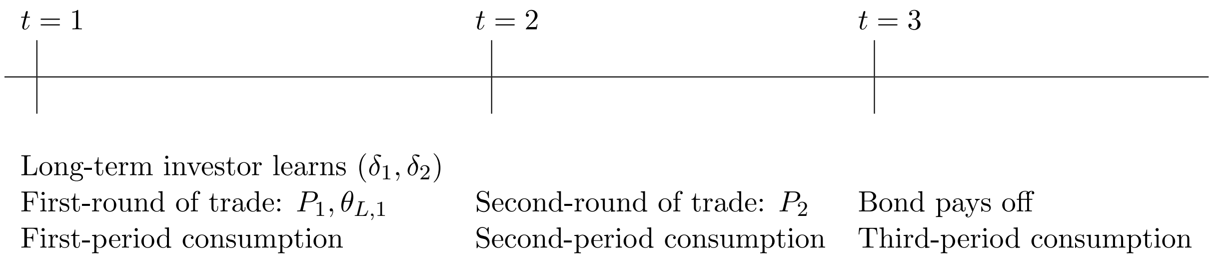

Figure 1: **Model Time-Line**

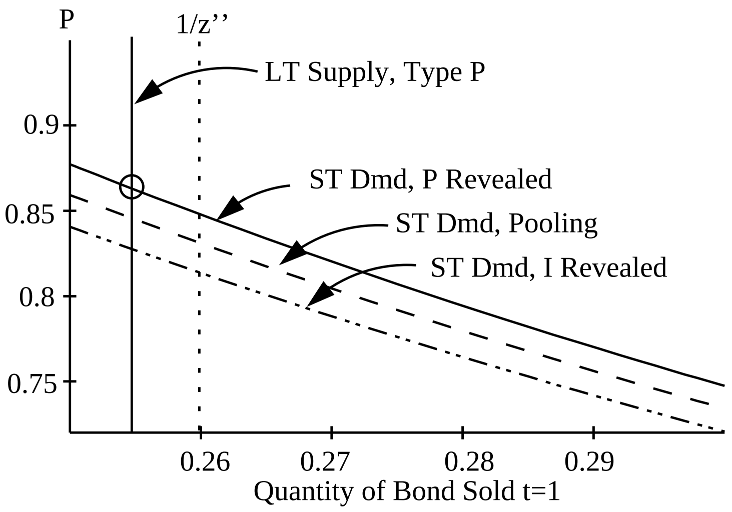

Figure 2a: **Example of a Fully Revealing Outcome — First Period.**

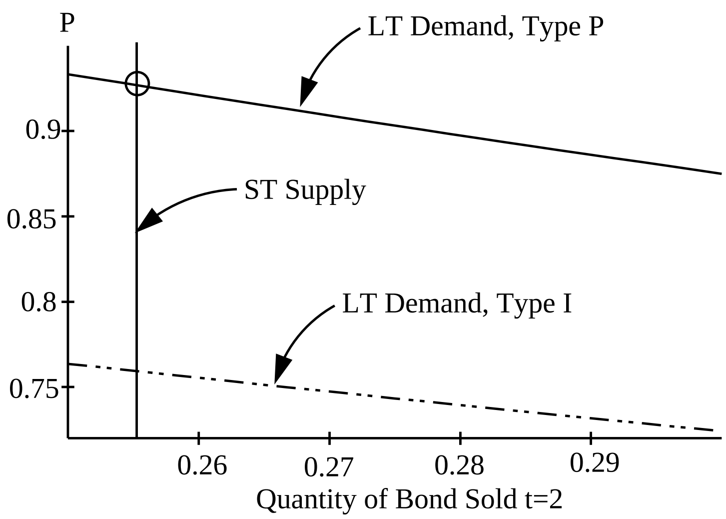

Figure 2b: **Example of a Fully Revealing Outcome — Second Period.**

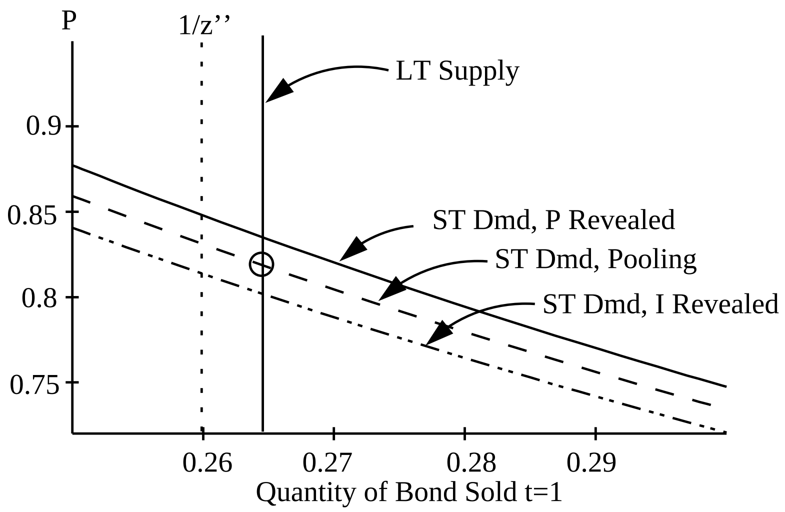

Figure 3a: **Example of a Partially Revealing Outcome — First Period.**

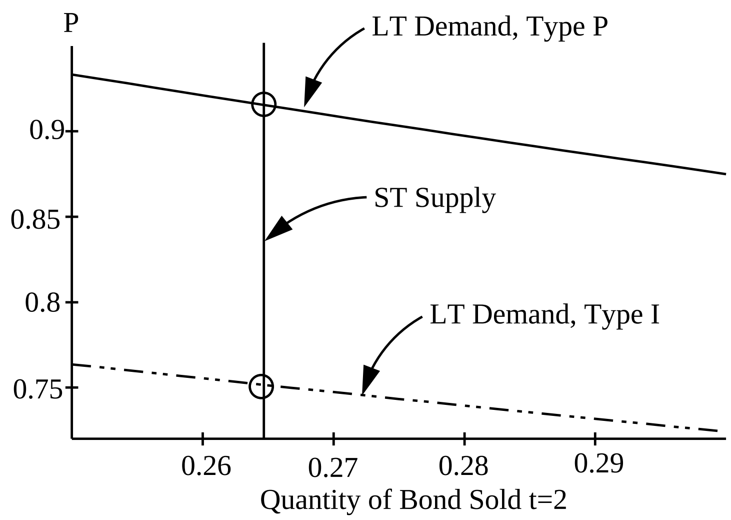

Figure 3b: **Example of a Partially Revealing Outcome — Second Period.**

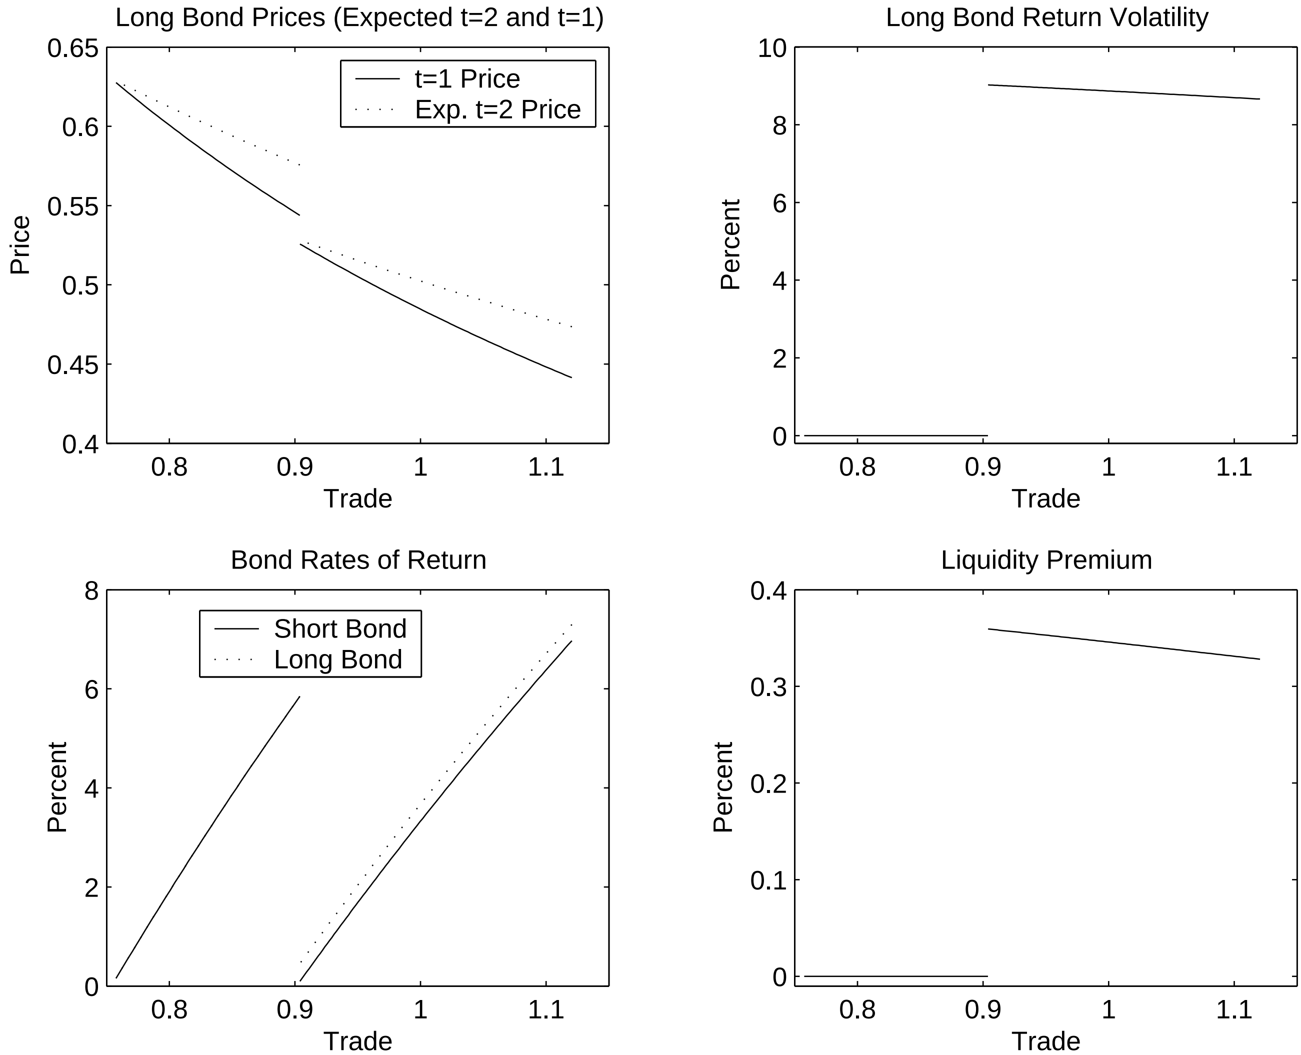

Figure 4: **Log Preferences and Uniform Priors Example.**

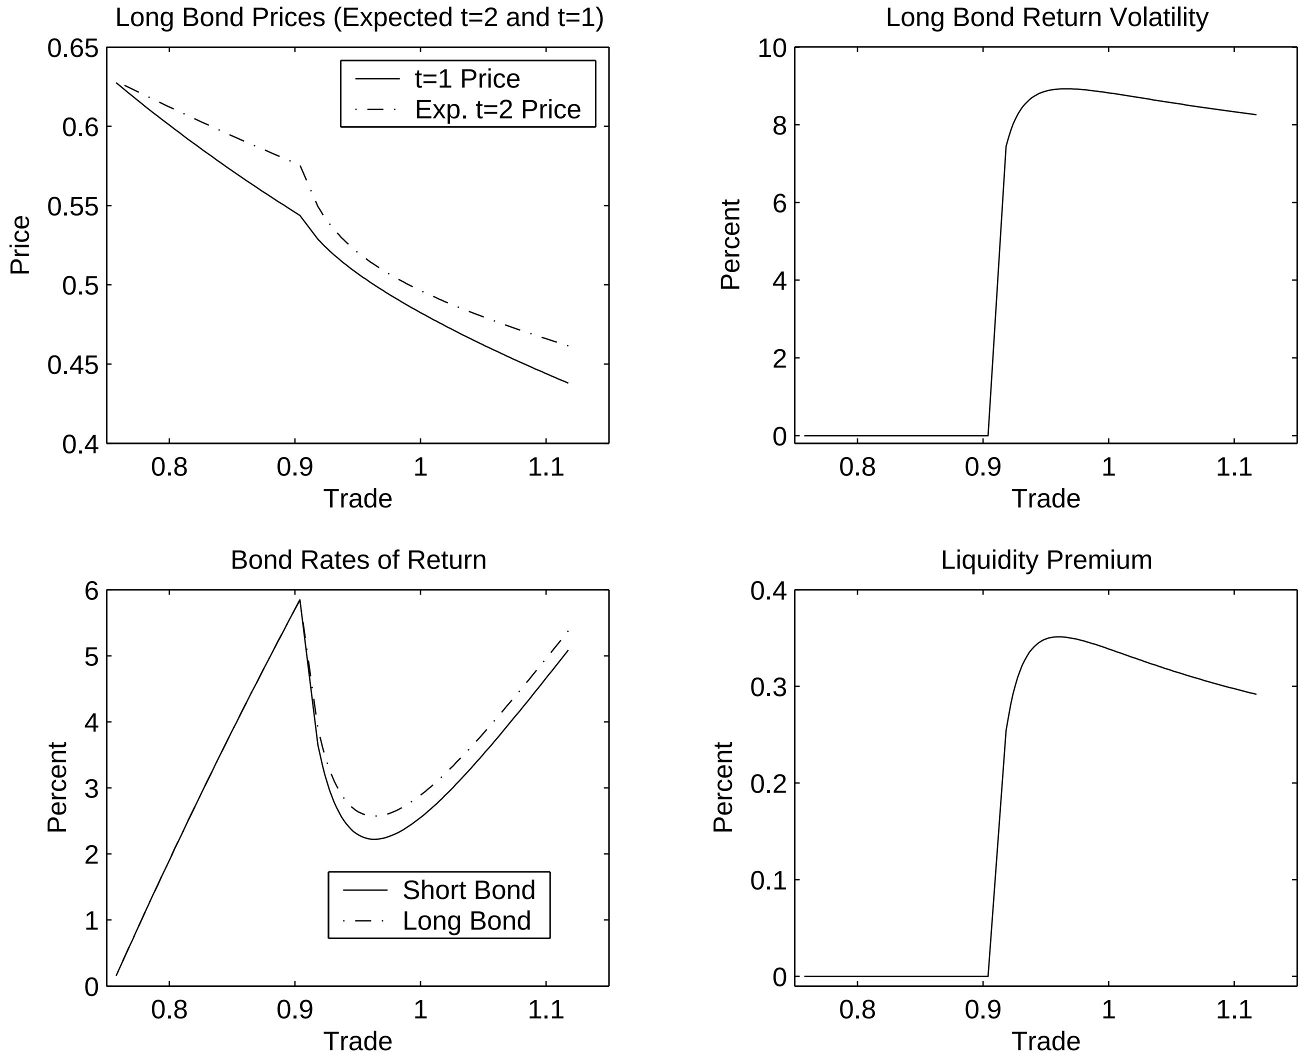

Figure 5: **Log Preferences and Beta Priors Example.**

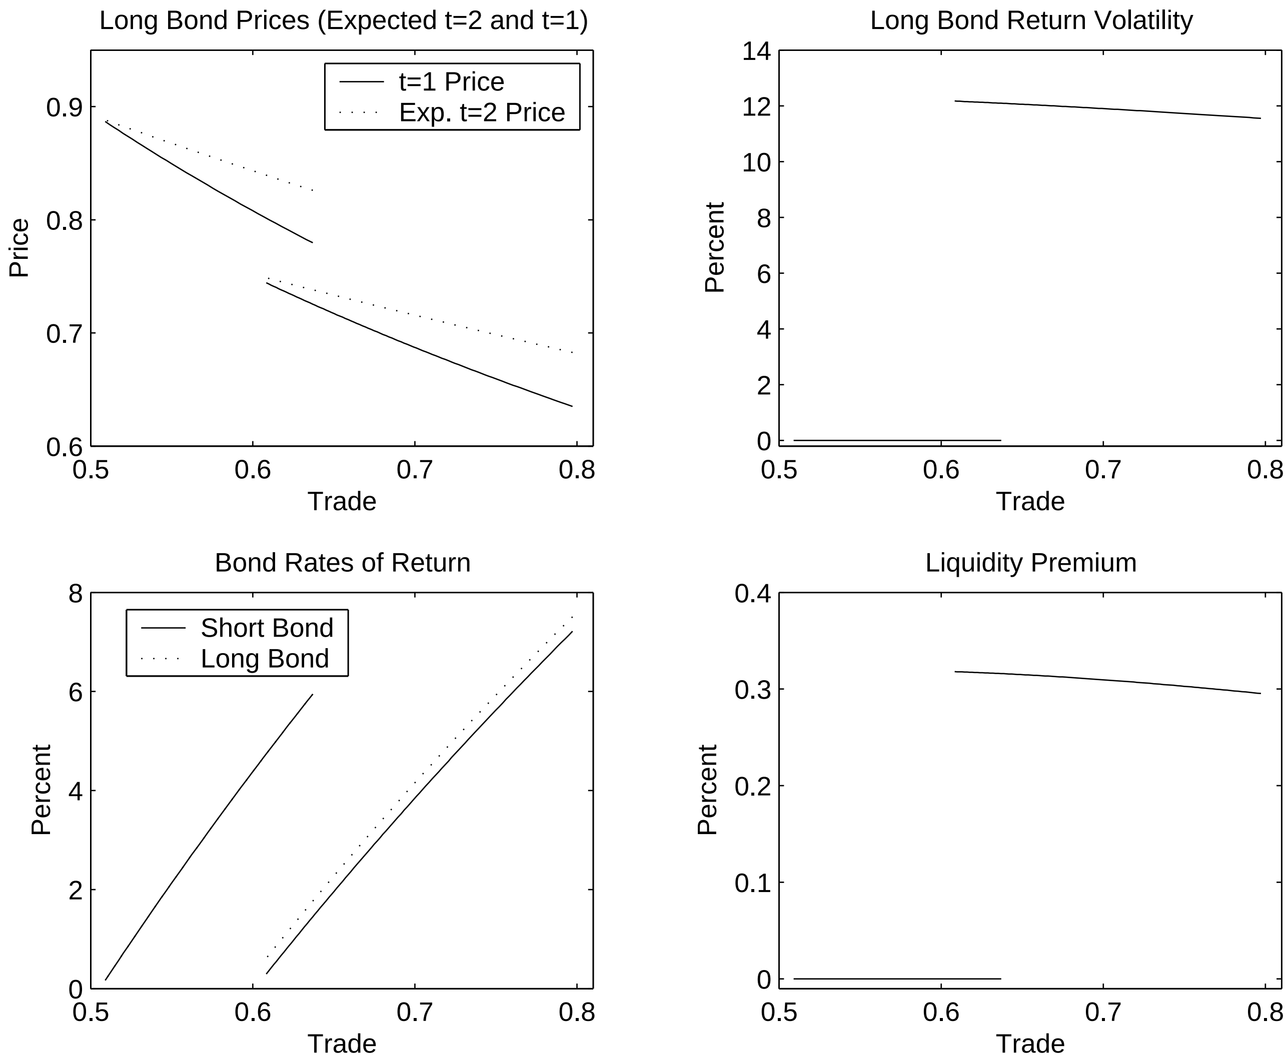

Figure 6: **Square Root Preferences and Uniform Priors Example.**

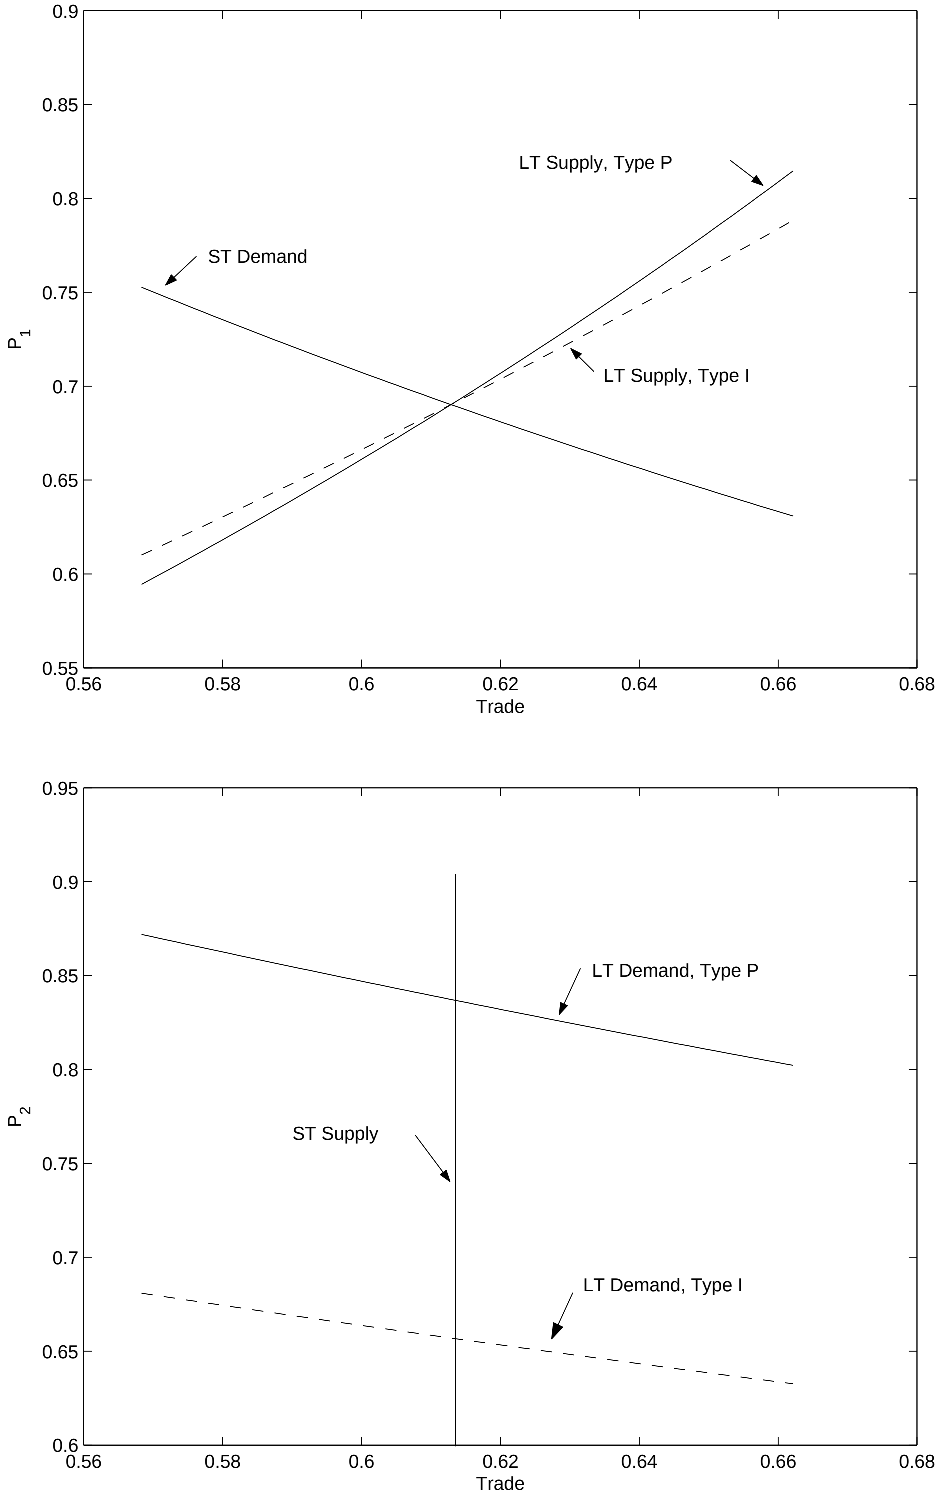

Figure 7: **Marshallian Demand Curves in the Square Root Preferences and Uniform Priors Example.**

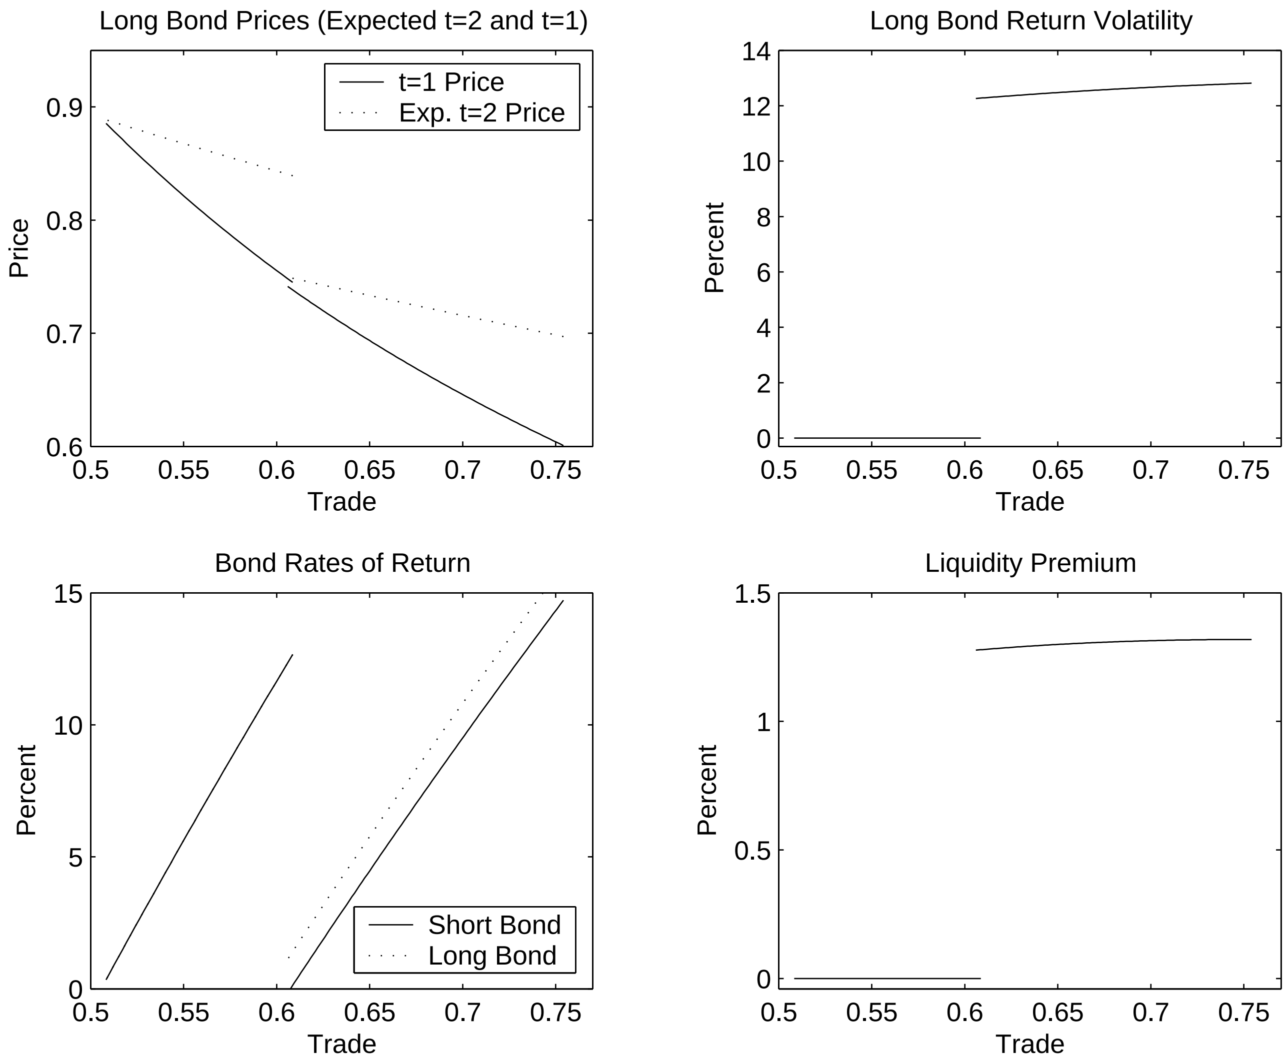

Figure 8: **Square Root Long-Horizon Preferences, CRRA=2 Short-Horizon Preferences, and Uniform Priors Example.**