

Ian Martin and Steve Ross

Notes on the yield curve

Article (Accepted version)
(Refereed)

Original citation: Martin, Ian and Ross, Steve (2018) *Notes on the yield curve*. [Journal of Financial Economics](#). ISSN 0304-405X

© 2018 [Elsevier B.V](#)

This version available at: <http://eprints.lse.ac.uk/90208/>
Available in LSE Research Online: September 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Notes on the yield curve

Ian Martin Steve Ross*

September, 2018

Abstract

We study the properties of the yield curve under the assumptions that (i) the fixed-income market is complete and (ii) the state vector that drives interest rates follows a finite discrete-time Markov chain. We focus in particular on the relationship between the behavior of the long end of the yield curve and the recovered time discount factor and marginal utilities of a pseudo-representative agent; and on the relationship between the “trappedness” of an economy and the convergence of yields at the long end.

JEL code: G12. Keywords: yield curve, term structure, recovery theorem, traps, Cheeger inequality, eigenvalue gap.

*Ian Martin (corresponding author): London School of Economics, London, UK, i.w.martin@lse.ac.uk. Steve Ross: MIT Sloan, Cambridge, MA. We are grateful to Dave Backus, Ravi Bansal, John Campbell, Mike Chernov, Darrell Duffie, Lars Hansen, Ken Singleton, Leifu Zhang, and to participants at the TIGER Forum and at the SITE conference for their comments. Ian Martin is grateful for support from the ERC under Starting Grant 639744.

In this paper, we present some theoretical results on the properties of the long end of the yield curve. Our results relate to two literatures. The first is the Recovery Theorem of Ross (2015). Ross showed that, given a matrix of Arrow–Debreu prices, it is possible to infer the objective state transition probabilities and marginal utilities. Although it is a familiar fact that sufficiently rich asset price data pins down the risk-neutral probabilities, it is initially surprising that we can do the same for the objective, or real-world, probabilities.

More precisely, the asset price matrix, \mathbf{A} , can be decomposed as $\mathbf{A} = \phi \mathbf{D} \mathbf{\Pi} \mathbf{D}^{-1}$, where $\phi > 0$ is a scalar, \mathbf{D} is a diagonal matrix with positive entries along the diagonal, and $\mathbf{\Pi}$ is a stochastic matrix (so its row sums all equal one). In Ross’s setting, with assets priced by an expected-utility-maximizing investor, ϕ can be interpreted as the investor’s time discount factor, and \mathbf{D} and $\mathbf{\Pi}$ summarize, respectively, the investor’s marginal utilities and subjective probabilities.

Here we emphasize that asset prices can be uniquely decomposed into objects that can be interpreted as probabilities and marginal utilities whether or not there is such an investor. We think of this result as establishing the existence of a *pseudo-representative agent*. We hypothesize that the recovered probabilities are the real-world probabilities. As Ross showed, this is true if there is a utility-maximizing investor—in which case this investor *is* the pseudo-representative agent—but more generally the question of whether the hypothesis holds is an empirical one. We discuss how to test it but do not settle the question here; instead we proceed on the assumption that the hypothesis holds and develop some of its implications.

The first is that recovery can be partially effected, without knowledge of the matrix \mathbf{A} , by observing the behavior of the long end of the yield curve. The yield on the (infinitely) long zero-coupon bond reveals the time preference rate of the pseudo-representative agent, and the time-series of returns on the long bond reveals the pseudo-representative agent’s marginal utilities. (These “marginal utilities” are a convenient guide to intuition, but our analysis is based on the logic of no-arbitrage:

we do not need to assume that the concept of expected utility is meaningful.)

These results connect the Recovery Theorem to earlier results of Backus, Gregory and Zin (1989), Kazemi (1992), Bansal and Lehmann (1994, 1997), Dybvig, Ingersoll and Ross (1996), Alvarez and Jermann (2005) and Hansen and Scheinkman (2009), though unlike most of these authors we focus exclusively on bonds and interest-rate derivatives rather than on the prices of claims to growing cashflows, for reasons discussed further below. A side benefit is that we can avoid some of the technical complication of this literature, and provide simple and short proofs of our results.

We derive several further results. First, we show that the yield curve must slope upwards on average. This is an empirical regularity. Second, we provide a formula for the expected excess return on the long bond in terms of the prices of options on the long bond. The matrix of Arrow–Debreu prices is in the background of both results, but we avoid the need to observe it directly. In Section 3, we define a notion of “trappedness” for an economy, and relate it to the speed of convergence (in yields and returns) at the long end of the yield curve. In Section 4, we address the question of when the matrix of Arrow–Debreu prices can be inferred from plausibly observable asset price data. We present several examples that illustrate the flexibility of our framework and exhibit some potential opportunities—and pitfalls—for empirical work.

1 The general framework

We work in discrete time. Our focus is on fixed-income assets, including fixed- or floating-rate bonds, interest-rate swaps, and derivatives on them, including caps, floors, bond options and swaptions. We assume that fixed-income markets are (i) complete, and (ii) governed by a state variable that follows a Markov chain on $\{1, 2, \dots, m\}$.

The state variable will in general be multidimensional (so that if the state vector

has J entries that can each take K values, the size of the state space is $m = K^J$. Although we have in mind that the state vector might in principle comprise, say, the current short rate, measures of the yield curve slope and curvature, the VIX index, state of the business cycle—even, perhaps, some measure of “animal spirits”—we will not need, in this paper, to be specific about its constituents. But as we assume that the state variable follows a Markov chain, it *is* important that all elements of the vector can reasonably be thought of as stationary.¹

In a complete market, it is convenient to summarize all available asset-price data in a matrix, \mathbf{A} , of Arrow–Debreu prices

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}(1, 1) & \mathbf{A}(1, 2) & \cdots & \mathbf{A}(1, m) \\ \mathbf{A}(2, 1) & \mathbf{A}(2, 2) & & \vdots \\ \vdots & & \ddots & \vdots \\ \mathbf{A}(m, 1) & \cdots & \cdots & \mathbf{A}(m, m) \end{pmatrix},$$

where we write $\mathbf{A}(i, j)$ for the price, in state i , of the Arrow–Debreu security that pays off \$1 if state j materializes next period. The absence of arbitrage requires that the entries of \mathbf{A} are nonnegative, and that $\mathbf{A}(i, j)$ is strictly positive if the transition from state i to state j occurs with positive probability. The row sums of \mathbf{A} are the prices of one-period riskless bonds in each state. We assume that not all of these row sums are equal, except where explicitly stated otherwise. This assumption is inessential, but it lets us avoid constantly having to qualify statements to account for the degenerate case in which interest rates are constant and the yield curve is flat.

We assume that for a sufficiently large time horizon, T , all Arrow–Debreu securities maturing in T periods have strictly positive prices in every state. This assump-

¹It has also long been accepted that fixed income yields are stationary. Much effort has been expended in statistical examination of this issue. But as Jon Ingersoll remarked during a seminar more than twenty-five years ago in which the speaker was using spectral analysis to show that interest rates were non-stationary, “Interest rates were about 5% 4000 years ago in Babylonian times and they’re still about 5% today—seems pretty stationary to me.” (John Cochrane and others have made the same point.)

tion, which we maintain throughout the paper, ensures that \mathbf{A} is a *primitive* matrix. Our framework is extremely flexible: it allows the yield curve to take any shape consistent with no arbitrage. For now, we think of the matrix \mathbf{A} as directly observable; we will address the issue that it may be hard to observe in practice below.

Example.—If there were a utility-maximizing investor with time discount factor ϕ , marginal utility $u'(i)$ in state i , and subjective transition probabilities $\pi(i, j)$, then the Arrow–Debreu prices would satisfy

$$\mathbf{A}(i, j) = \phi\pi(i, j)\frac{u'(j)}{u'(i)}. \quad (1)$$

In terms of inverse marginal utility, $v(i) \equiv 1/u'(i)$, we would have $\mathbf{A}(i, j) = \phi v(i)\pi(i, j)/v(j)$ for each i and j or, written more concisely as a matrix equation, $\mathbf{A} = \phi\mathbf{D}\mathbf{\Pi}\mathbf{D}^{-1}$, where \mathbf{D} is a diagonal matrix with positive diagonal elements $\{v(i)\}$ and $\mathbf{\Pi} = \{\pi(i, j)\}$ is a *stochastic* matrix, i.e. a matrix whose row sums all equal 1 (because they sum over the probabilities of moving from the current state to any other state).

The next result shows quite generally that asset price data, \mathbf{A} , can be decomposed into a diagonal matrix \mathbf{D} with positive diagonal entries and a stochastic matrix $\mathbf{\Pi}$. Since the decomposition exists whether or not there exists a utility-maximizing investor, we prefer to think of this result as establishing the existence of a *pseudo-representative agent*.

Result 1. *Given asset price data represented by the primitive matrix \mathbf{A} , there exists a unique decomposition*

$$\mathbf{A} = \phi\mathbf{D}\mathbf{\Pi}\mathbf{D}^{-1} \quad (2)$$

where \mathbf{D} is a diagonal matrix with positive entries along the diagonal and $\mathbf{\Pi}$ is a transition matrix (up to economically irrelevant scalar multiples of \mathbf{D}).

Proof. The Perron–Frobenius theorem states that \mathbf{A} has a unique (up to scale) eigenvector, \mathbf{v} , and real eigenvalue, ϕ , satisfying $\mathbf{A}\mathbf{v} = \phi\mathbf{v}$, where ϕ and \mathbf{v} are strictly

positive and ϕ is the largest, in absolute value, of all the eigenvalues of \mathbf{A} . Letting \mathbf{D} be the diagonal matrix with \mathbf{v} along its diagonal and \mathbf{e} be the vector of ones, so that $\mathbf{v} = \mathbf{D}\mathbf{e}$,

$$\mathbf{\Pi} \equiv \frac{1}{\phi} \mathbf{D}^{-1} \mathbf{A} \mathbf{D}$$

is a transition matrix, i.e., it is positive and $\mathbf{\Pi}\mathbf{e} = \mathbf{e}$.

Now suppose that there is a different decomposition of \mathbf{A} as in (2) with a diagonal matrix \mathbf{C} . By the decomposition it follows that $\mathbf{C}\mathbf{e}$ is an eigenvector of \mathbf{A} , and as by the Perron–Frobenius theorem \mathbf{A} has a unique positive eigenvector it follows that $\mathbf{C}\mathbf{e}$ is proportional to $\mathbf{D}\mathbf{e}$, and therefore that (up to scale) $\mathbf{C} = \mathbf{D}$, with ϕ as the unique eigenvalue. It is immediate that $\mathbf{\Pi}$ is also unique. \square

Result 1 is a generalization of the Recovery Theorem of Ross (2015); if $\mathbf{\Pi}$ is the true probability transition matrix and the stochastic pricing kernel is (in Ross’s terminology) transition-independent then Result 1 assures that $\mathbf{\Pi}$ can be uniquely recovered from \mathbf{A} . But the decomposition is well defined and unique whether or not there is a representative agent and there is no need to assume that there is a representative agent. Rather, this result can be thought of as establishing the existence and uniqueness of a pseudo-representative agent with marginal utilities and time discount factor encapsulated in \mathbf{D} and ϕ , respectively. We will take the following hypothesis as a given.

Hypothesis 1. *The probability distribution $\mathbf{\Pi}$ recovered in the above fashion is the true objective probability distribution.*

Whether this hypothesis holds is fundamentally an empirical question and is not an issue to be determined on *a priori* theoretical grounds; it can only be resolved by comparing the recovered distribution, $\mathbf{\Pi}$, with observed realizations in the market. In much of this paper, we will assume that the hypothesis holds, though in Result 5 we provide an interpretation of the probabilities $\mathbf{\Pi}$ that is valid even if Hypothesis 1 is false. There is a parallel here with the expectations hypothesis, one version of

which asserts that bond pricing is risk-neutral. This has provided a useful way to think about interest rates and related phenomena and is the basis for a substantial literature.

We have now linked the eigenvector, \mathbf{v} , and eigenvalue ϕ to, respectively, the marginal utilities and time discount factor of the pseudo-representative agent. On the face of it, we are still confronted with the difficult task of having to compute the matrix \mathbf{A} in order to find its eigenvalues and eigenvectors. It turns out, though, that we can avoid this difficulty by connecting \mathbf{v} and ϕ to the long end of the yield curve.

2 The yield curve

Let $B_t(i)$ denote the price of a zero-coupon bond paying 1 in t periods if the current (time zero) state is i . The continuously-compounded t -period yield, $y_t(i)$, and the simple yield, $Y_t(i)$, are given by

$$B_t(i) = e^{-y_t(i)t} = \frac{1}{[1 + Y_t(i)]^t}.$$

By the law of iterated expectations, we have

$$B_t(i) = \sum_j \mathbf{A}^t(i, j), \tag{3}$$

where $\mathbf{A}^t(i, j)$ is the entry that appears in the i th row and j th column of the matrix $\mathbf{A}^t \equiv \underbrace{\mathbf{A} \cdots \mathbf{A}}_{t \text{ times}}$. In other words, to calculate the price in state i of a t -period zero-coupon bond, we sum the elements of the i th row of \mathbf{A}^t . Another way to see this is to observe that (given our Markov chain assumption) element (i, j) of \mathbf{A}^t is the current price, in state i , of a pure contingent security that pays 1 in t periods if the state is then j ; and the price of a riskless bond is the sum, over all states j , of the prices of these contingent securities.

It follows that

$$y_t(i) = -\frac{1}{t} \log \sum_j \mathbf{A}^t(i, j)$$

and the *short rate*, in state i , is

$$r_f(i) = y_1(i) = -\log \sum_j \mathbf{A}(i, j).$$

The *long rate* in state i is defined as

$$y_\infty(i) = -\lim_{t \rightarrow \infty} \frac{1}{t} \log \sum_j \mathbf{A}^t(i, j).$$

The realized return on a t -period bond from a transition from state i to state j in a single period is

$$R_t(i, j) \equiv \frac{B_{t-1}(j)}{B_t(i)},$$

and $R_\infty(i, j) \equiv \lim_{t \rightarrow \infty} R_t(i, j)$. We use lowercase letters for log returns,

$$r_t(i, j) \equiv \log R_t(i, j) \quad \text{and} \quad r_\infty(i, j) \equiv \log R_\infty(i, j).$$

We obtain expected returns by summing over transition probabilities. For example, the conditionally expected return on the long bond in state i is

$$\bar{R}_\infty(i) \equiv \sum_j \pi(i, j) R_\infty(i, j)$$

and the unconditional expected log return on the long bond is

$$\bar{r}_\infty \equiv \sum_{i, j} \pi(i) \pi(i, j) r_\infty(i, j). \tag{4}$$

The quantity $\pi(i)$ that appears in (4) is the long-run probability that the economy is in state i —that is, the i th element of the stationary distribution $\boldsymbol{\pi}$, which is the left

eigenvector of $\mathbf{\Pi}$ that satisfies

$$\boldsymbol{\pi}'\mathbf{\Pi} = \boldsymbol{\pi}'. \quad (5)$$

While the focus above has been on \mathbf{v} , a right-eigenvector of \mathbf{A} , the corresponding left-eigenvector \mathbf{w} , which satisfies $\mathbf{w}'\mathbf{A} = \phi\mathbf{w}'$, also has an interesting interpretation. It follows from (2) and (5) that (up to a scalar multiple) $\mathbf{w} = \mathbf{D}^{-1}\boldsymbol{\pi}$, and hence that $w(i) = \pi(i)/v(i)$. We will see below that this can be interpreted as (proportional to) the risk-neutral stationary distribution. An important limiting result (see Theorem 8.5.1 of Horn and Johnson (1990)) is that

$$\lim_{t \rightarrow \infty} \frac{\mathbf{A}^t}{\phi^t} = \mathbf{v}\mathbf{w}' = [v(i)w(j)], \quad (6)$$

where \mathbf{v} and \mathbf{w} are normalized so that $\mathbf{w}'\mathbf{v} = 1$. This fact is exploited in the next result to give the eigenvector \mathbf{v} another economic interpretation.

Result 2. *Let the \mathbf{v} -asset be the asset that pays off $v(j)$ next period. Returns on the \mathbf{v} -asset are identical to the returns on the long bond.*

Proof. The price of the \mathbf{v} -asset is equal to

$$\sum_j \mathbf{A}(i, j)v(j) = \phi v(i),$$

using the fact that \mathbf{v} is an eigenvector of \mathbf{A} . So the return on the \mathbf{v} -asset is $v(j)/[\phi v(i)]$. Using (6), the return on the long bond is

$$R_\infty(i, j) = \lim_{t \rightarrow \infty} \frac{\sum_k \mathbf{A}^{t-1}(j, k)}{\sum_k \mathbf{A}^t(i, k)} = \frac{1}{\phi} \lim_{t \rightarrow \infty} \frac{\sum_k \frac{\mathbf{A}^{t-1}(j, k)}{\phi^{t-1}}}{\sum_k \frac{\mathbf{A}^t(i, k)}{\phi^t}} = \frac{\sum_k v(j)w(k)}{\phi \sum_k v(i)w(k)} = \frac{v(j)}{\phi v(i)}. \quad (7)$$

□

Empirically, then, if we could observe the return on the long bond Result 2 would

allow us to observe the kernel as well. (Kazemi (1992) was the first to observe a connection of this type; he did so in a continuous-time diffusion model.) Below we will see whether it is possible to approximate this return from the returns on sufficiently long dated bonds. We now link the eigenvalue ϕ to the long end of the yield curve.

Result 3. *The long rate and unconditional expected log return on the long bond both equal $-\log \phi$.*

Proof. We have

$$\begin{aligned}
y_\infty(i) &= -\lim_{t \rightarrow \infty} \frac{1}{t} \log \sum_j \mathbf{A}^t(i, j) \\
&= -\lim_{t \rightarrow \infty} \frac{1}{t} \log \left[\phi^t \sum_j \frac{\mathbf{A}^t}{\phi^t}(i, j) \right] \\
&= -\log \phi - \lim_{t \rightarrow \infty} \frac{1}{t} \log \sum_j \frac{\mathbf{A}^t}{\phi^t}(i, j) \\
&= -\log \phi,
\end{aligned}$$

where the final equality holds because, as noted in equation (6), \mathbf{A}^t/ϕ^t tends to a constant matrix.

From (7) we have $r_\infty(i, j) = -\log \phi + \log[v(j)/v(i)]$, so

$$\bar{r}_\infty = -\log \phi + \sum_{i,j} \pi(i)\pi(i, j) \log \frac{v(j)}{v(i)} = -\log \phi + \sum_j \pi(j) \log v(j) - \sum_i \pi(i) \log v(i) = -\log \phi,$$

where we have exploited the fact that $\pi' \mathbf{\Pi} = \pi'$. □

Given the result of Dybvig, Ingersoll and Ross (1996), it is no surprise that the long yield is a constant. Nonetheless, it might at first seem contradictory to have the long rate converge to a constant while the return on the long bond is variable. But

the return on a t -period bond is

$$r_t(i, j) = \underbrace{y_t(i)}_{\text{yield}} - \underbrace{(t-1)}_{\text{duration}} \times \underbrace{(y_{t-1}(j) - y_t(i))}_{\text{realized yield change}}.$$

While Result 3 shows that the realized yield change is very small for long-maturity bonds (and zero in the limit), the duration, $t - 1$, is very large (and infinite for the limiting long bond). When the two effects are multiplied, the result is that long-dated bonds have volatile returns.

Summarizing the results so far, returns on the long bond reveal the eigenvector \mathbf{v} , i.e., the pricing kernel, and the long yield reveals the eigenvalue ϕ , and hence the time discount factor of the pseudo-representative agent.

We now turn to the question of making inferences about the objective probability distribution from observable data. As a crude example of the kind of thing we are interested in, note that Result 3 implies (by Jensen's inequality) that the expected return on the long bond is at least as great as its simple yield, $\bar{R}_\infty > 1 + Y_\infty$. We will shortly improve on this by deriving a formula that expresses the conditionally expected return on the long bond in terms of option prices. The next result is a stepping stone towards that goal.

Result 4. *Within the class of fixed-income assets, the long bond is growth-optimal; that is, in every state the long bond has the highest expected log return of all fixed-income assets. Hence the long bond earns a positive risk premium in every state (other than in states that evolve deterministically to a fixed successor state, in which case the long bond is riskless).*

Proof. The state- i price of an asset with payoffs \mathbf{x} is

$$\begin{aligned} \sum_j \mathbf{A}(i, j)x(j) &\stackrel{(a)}{=} \phi \sum_j \pi(i, j) \frac{v(i)}{v(j)} x(j) \\ &\stackrel{(b)}{=} \sum_j \pi(i, j) \frac{x(j)}{R_\infty(i, j)} \end{aligned} \tag{8}$$

Equality (a) follows from Result 1, and equality (b) follows from Result 2. It follows that the reciprocal of the return on the long bond is a stochastic discount factor. But then the long bond's return must be growth-optimal, as in a complete market the reciprocal of the growth-optimal return is the unique stochastic discount factor.

The final statement follows because the long bond's log expected return is weakly greater than its expected log return, which is weakly greater than the log riskless rate (by Jensen's inequality and the growth-optimality of the long bond, in turn). If the state evolution is nondeterministic, the equalities are strict. \square

Equation (8) shows that a fixed-income payoff can be priced either via state prices—equivalently, by discounting at the short riskless rate and using risk-neutral probabilities—as on the left-hand side, or by discounting at the long bond's return and using the true probabilities, as on the right-hand side. Loosely speaking, Result 4 can be thought of as saying that the cheapest way of generating a payoff in the far-distant future is to buy the growth-optimal portfolio, or equivalently the long bond.

A variant of the above result provides a more general interpretation of $\mathbf{\Pi}$ that does not rely on Hypothesis 1.

Result 5. *Whether or not Hypothesis 1 holds, the entries $\pi(i, j)$ of $\mathbf{\Pi}$ represent the risk-neutral probabilities with the long bond as numeraire.² It follows that the recovered probabilities $\mathbf{\Pi}$ can be interpreted as the true probabilities perceived by an agent with log utility who chooses to invest his or her wealth fully in the long bond (or, equivalently, in the \mathbf{v} -asset).*

Proof. The characterization of $\mathbf{\Pi}$ as the risk-neutral probabilities with the long bond as numeraire follows immediately from equation (8). The second statement follows

²Probabilities $\hat{\pi}(i, j)$ are risk-neutral using the long bond as numeraire if any payoff \mathbf{x} has price, in state i , equal to $\sum_j \hat{\pi}(i, j) \frac{\mathbf{x}(j)}{R_\infty(i, j)}$. (For comparison, the conventional risk-neutral probabilities $\pi^*(i, j)$ are risk-neutral *using the short riskless bond as numeraire*, so the price can also be expressed as $\sum_j \pi^*(i, j) \frac{\mathbf{x}(j)}{R_f(i)}$.) The notion of risk-neutral probabilities defined relative to numeraires was first explicitly introduced by Geman, El Karoui and Rochet (1995).

because such an investor’s stochastic discount factor is $1/R_\infty$. □

Result 5 clarifies the implicit assumption that underlies our (and Ross’s (2015)) hypothesis that $\mathbf{\Pi}$ is the true probability transition matrix. The perspective of a log investor was also adopted at various points in Martin (2017) and Kremens and Martin (2018); but in those papers, the log investor was assumed to be fully invested in the *stock market*. Given the work of Alvarez and Jermann (2005), who argue that the average log return on the stock market is larger than that of the long bond, it is doubtful that the perspective of the log investor who is fully invested in bonds—as in Result 5—is consistent with equity pricing.

To see this another way, consider the matrix, \mathbf{A} , used by Ross (2015, Table II, Panel B) to illustrate recovery in the equity market. There are 11 states in the example. The yield curves implied by \mathbf{A} , in each of the states, are shown in Figure 1. (By decomposing \mathbf{A} into matrices \mathbf{D} and $\mathbf{\Pi}$ as in Result 1, one can also check that the economy spends more than 90% of its time in the three most extreme states, one with short rate below -8% and two with short rates above 10% .) The figure provides an alternative way to understand the points made by Alvarez and Jermann (2005) and by Borovička, Hansen and Scheinkman (2017). In this framework, pricing is risk-neutral if the riskless rate is constant, as was shown in Ross (2015). Conversely, confronted with the large risk premium available in the equity market, the framework is forced to conclude (counterfactually) that the riskless rate fluctuates wildly.

Whether or not the long bond prices all assets including equities, though, there is nothing in this evidence that prevents it from being the projection of the economy-wide stochastic discount factor on the fixed income market. In that sense, as there is agreement that yields are stationary, we are comfortable of thinking of it as the appropriate discount factor for stationary assets.

While the returns on the long bond replicate the pricing kernel, there are other assets that also can serve as the pricing operator. The next result shows that the (unique) infinitely-lived asset with a constant dividend yield is also a surrogate for




Figure 1: Yield curves in each of the 11 states in an example estimated by Ross (2015).

the stochastic discount factor. To see this, consider an infinitely lived asset that pays, in every period, $x(i)$ if the economy is in state i , for all i . (If $x(i)$ is constant across i then the asset is simply a consol.) The value, p_∞ , of the asset is

$$p_\infty = \mathbf{A}\mathbf{x} + \mathbf{A}^2\mathbf{x} + \dots = \mathbf{A}^*\mathbf{x},$$

where

$$\mathbf{A}^* \equiv \mathbf{A} + \mathbf{A}^2 + \dots = \mathbf{A}(\mathbf{I} - \mathbf{A})^{-1}$$

converges because $\phi < 1$ by assumption. Notice that \mathbf{A}^* inherits the same dominant eigenvector as \mathbf{A} , namely \mathbf{v} , and that the associated maximal eigenvalue is $\frac{\phi}{1-\phi}$.

Result 6. *There is a unique infinitely-lived, limited liability asset with a constant dividend yield. Its dividend yield is $D/P_\infty \equiv 1/\phi - 1$ and its returns perfectly replicate the returns on the long bond and on the one-period \mathbf{v} -asset. No asset can have a uniformly higher or lower dividend yield than this asset.*

Proof. Consider the perpetual \mathbf{v} asset, which pays \mathbf{v} every period. The price of this asset is $\mathbf{A}^*\mathbf{v} = \frac{\phi}{1-\phi}\mathbf{v}$, so its dividend yield is constant at $1/\phi - 1 = Y_\infty$. Uniqueness follows because any asset with constant dividend yield is an eigenvector of \mathbf{A}^* ; and

we have seen that up to multiples there is a unique eigenvector that is positive (as required for limited liability). The return on this asset, on moving from state i to state j , is $[v(j) + \frac{\phi}{1-\phi}v(j)] / [\frac{\phi}{1-\phi}v(i)] = v(j) / [\phi v(i)]$. The last claim follows by applying Theorem 8.1.26 of Horn and Johnson (1990) to \mathbf{A}^* . \square

We emphasize that the asset with constant dividend yield is not a perpetuity except in the special case in which interest rates are constant.

The above results allow us to uncover some new results on pricing in fixed income markets. We define the t -period forward rate $f_t(i) = -\log B_t(i) + \log B_{t-1}(i)$.

Result 7. *On average, the forward curve lies below the long yield.*

Proof. Notice that $f_t(i) = r_t(i, j) + \log B_{t-1}(i) - \log B_{t-1}(j)$. It follows—using the fact that $\sum_{i,j} \pi(i)\pi(i, j)(\log B_{t-1}(i) - \log B_{t-1}(j)) = 0$, and Results 3 and 4—that $\sum_i \pi(i)f_t(i) = \sum_{i,j} \pi(i)\pi(i, j)f_t(i) = \sum_{i,j} \pi(i)\pi(i, j)r_t(i, j) < \bar{r}_\infty = y_\infty$. \square

As an immediate corollary, we have

Result 8. *On average, the yield curve lies below the long yield.³ (But the yield curve cannot always lie below the long yield.)*

Proof. The first claim follows from Result 7, because $y_t(i) = \frac{-1}{t} \log B_t(i) = \frac{1}{t} \sum_{s=1}^t f_s(i)$. The second follows because $\min_i \sum_j \mathbf{A}^t(i, j) \leq \phi^t$ (see, for example, Theorem 8.1.22 of Horn and Johnson (1990)), and hence $\max_i y_t(i) \geq y_\infty$. \square

Together with our earlier results, Result 8 implies that the time discount factor $\phi < e^{-\bar{y}_t}$ for any t : average yields provide an upper bound on the subjective time discount factor of the pseudo-representative agent. This strengthens the finding in Ross (2015) that the maximal short rate provides an upper bound on the time discount factor.

³It is tempting to conjecture that the approach to the long run is monotone in Results 7 and 8—for example, that the average t -period yield \bar{y}_t or forward rate \bar{f}_t is increasing in t . But this is not true in general.

We have already seen that the yield and expected log return of the long bond are constant. But this does *not* imply that the expected arithmetic return on the long bond is constant, as long-dated bonds have volatile returns due to their large durations, which scale up the influence of tiny fluctuations in the yield curve. The next result shows how conditional moments of the return on the long bond can in principle be determined from option prices. It adapts the results of Martin (2017) and Martin and Wagner (2018) (relating expected returns on the stock market or on individual stocks to the prices of index or individual stock options) to the case of the long bond, exploiting the fact that the long bond is growth-optimal in our framework.

Result 9. *Options on the long bond reveal its conditional expected excess return:*

$$\overline{R}_\infty(i) - R_f(i) = 2 \left\{ \int_0^{R_f(i)} \text{put}(K; i) dK + \int_{R_f(i)}^\infty \text{call}(K; i) dK \right\}, \quad (9)$$

where $\text{call}(K; i)$ is the price, in state i , of a call option with strike K , maturing next period, on the long bond return, and $\text{put}(K; i)$ is the corresponding put price.

More generally, option prices reveal all the conditional moments of the return on the long bond. The n th conditional moment of the long bond return, $\overline{R}_\infty^n(i)$, satisfies

$$\overline{R}_\infty^n(i) - R_f(i)^n = n(n+1) \left\{ \int_0^{R_f(i)} K^{n-1} \text{put}(K; i) dK + \int_{R_f(i)}^\infty K^{n-1} \text{call}(K; i) dK \right\}.$$

Proof. The first statement is a special case of the second, which we now prove. Suppose we are in state i . Substitute $x(j) = R_\infty(i, j)^{n+1}$ in equation (8):

$$\sum_j \mathbf{A}(i, j) R_\infty(i, j)^{n+1} = \sum_j \pi(i, j) R_\infty(i, j)^n. \quad (10)$$

The right-hand side is the desired conditional moment, $\overline{R}_\infty^n(i)$. The left-hand side is the price of a claim to the $(n+1)$ -th power of the long bond return, settled next period. If options on the long bond return are traded, this payoff can be priced by a static no-

arbitrage argument. To do so, note that $x^{n+1} = n(n+1) \int_0^\infty K^{n-1} \max\{0, x - K\} dK$ for arbitrary $x \geq 0$. Setting $x = R_\infty(i, j)$, multiplying on both sides by $\mathbf{A}(i, j)$, summing over j , and interchanging sum and integral, this implies that

$$\sum_j \mathbf{A}(i, j) R_\infty(i, j)^{n+1} = n(n+1) \int_0^\infty K^{n-1} \underbrace{\sum_j \mathbf{A}(i, j) \max\{0, R_\infty(i, j) - K\}}_{\text{call}(K; i)} dK.$$

The result follows by splitting the range of integration and using the put–call parity relation $\text{call}(K; i) - \text{put}(K; i) = 1 - K/R_f(i)$. \square

Unfortunately it is difficult to test the above result directly because options on the long bond are not observable in practice. We therefore view it as indicative of a direction that empirical work might take.⁴ The question then becomes: how fast do returns on bonds of long-but-finite maturity approach the returns on the long bond? We now turn to this issue.

3 Traps and convergence at the long end

We have seen that the long end of the yield curve appropriately defined converges to the unknown pricing kernel, $v(i)$. In this section we will explore the speed of this convergence. We start by introducing the metric

$$Q \equiv \log \frac{\max_k v(k)}{\min_k v(k)} \geq 0.$$

⁴Bakshi, Chabi-Yo and Gao (2018) attempt to test the prediction (9) of Result 9 using what they describe as “options on the 30-year Treasury bond futures.” This characterisation is potentially misleading, however, as (i) the deliverables for a Treasury bond futures contract are bonds with maturities between 15 and 25 years, and (ii) these are *coupon* bonds, so have durations shorter than their maturities: a 15-year bond trading at par with a 5% coupon has a modified duration of about 10 years. Their analysis implicitly assumes that the returns on such bonds accurately reflect the returns on an infinite-duration bond; it also neglects the cheap-to-deliver option, which gives the futures contract negative convexity.

We can loosely think of Q as measuring the extent to which risk aversion matters for pricing. If pricing is risk-neutral then the pricing kernel $\mathbf{v} = \mathbf{e}$ and $Q = 0$. Put another way, if $Q = 0$ then \mathbf{e} is an eigenvector of \mathbf{A} which implies that interest rates are constant. This means that the yield curve is flat so that pricing (of fixed-income securities) is risk-neutral and the risk-neutral and objective probabilities coincide. From Result 2 we can see that Q is also a measure of the dispersion of long bond returns, as from equation (7) we have $\min_{i,j} r_\infty(i, j) = y_\infty - Q$ and $\max_{i,j} r_\infty(i, j) = y_\infty + Q$. This allows us to prove the following convergence result.

Result 10. *Q bounds the difference between the t -period yield and the long yield:*

$$|y_t(i) - y_\infty| \leq \frac{Q}{t}.$$

It therefore bounds the difference in yields of any two bonds of maturities t_1 and t_2 ,

$$|y_{t_1}(i) - y_{t_2}(i)| \leq Q \left(\frac{1}{t_1} + \frac{1}{t_2} \right).$$

Fixing the maturity t , we can bound the change in the yield uniformly across states,

$$|y_t(i) - y_t(j)| \leq \frac{2Q}{t}.$$

Proof. The price of a t -period bond in state i is $B_t(i) = \sum_j \mathbf{A}^t(i, j)$. Now, for all i ,

$$\phi^t \min_k v(k) \leq \phi^t v(i) = \sum_j \mathbf{A}^t(i, j) v(j) \leq \max_k v(k) \sum_j \mathbf{A}^t(i, j),$$

and similarly

$$\phi^t \max_k v(k) \geq \phi^t v(i) = \sum_j \mathbf{A}^t(i, j) v(j) \geq \min_k v(k) \sum_j \mathbf{A}^t(i, j).$$

It follows that $\phi^t e^{-Q} \leq B_t(i) \leq \phi^t e^Q$ for all i , and hence that

$$y_\infty - \frac{Q}{t} \leq y_t(i) \leq y_\infty + \frac{Q}{t},$$

which establishes all three results. \square

To summarize, Q controls the rate at which finite-maturity yields approach the long yield. If Q is small, then the yield curve must be fairly flat and yield volatility low. Conversely, if there is substantial variation in yields either across *maturities* or across *states of the world*—if the yield curve has significant slope, or if yields are volatile—then Q is large and risk considerations are important for fixed-income pricing.

The next result provides an analogous bound on the *returns* of long-dated zero-coupon bonds and, indeed, on any long-dated asset with a single payoff at time T .

Result 11. *The return on a long-dated asset paying $x(j)$ in state j at time T (and zero for $t \neq T$) approaches the return on the long bond as $T \rightarrow \infty$. More precisely,*

$$R_T(i, j) = R_\infty(i, j) + O(\delta^T), \tag{11}$$

where $\psi/\phi < \delta < 1$, ψ is the second-largest of the absolute values of the eigenvalues of \mathbf{A} , and δ can be chosen arbitrarily close to ψ/ϕ . Thus the eigenvalue gap ψ/ϕ determines how rapidly long-dated assets' returns converge to the return on the long bond.

Proof. The price of the asset in state i is $(\mathbf{A}^T \mathbf{x})_i$. We can conclude from Theorem 8.5.1 of Horn and Johnson (1990) that $\mathbf{A}^T(i, j) = \phi^T v(i)w(j) + O(\zeta^T)$ where ζ , which satisfies $\psi < \zeta < \phi$, can be chosen to be arbitrarily close to ψ . Therefore the asset's

price in state i is

$$\text{price} = \underbrace{\phi^T}_{\text{time}} \times \underbrace{v(i)}_{\text{state}} \times \underbrace{\sum_j w(j)x(j)}_{\text{asset-specific risk}} + O(\zeta^T).$$

Writing $K = \sum_j w(j)x(j)$ for the asset-specific risk term, the asset's return equals

$$R_T(i, j) = \frac{K\phi^{T-1}v(j) + O(\zeta^{T-1})}{K\phi^T v(i) + O(\zeta^T)}.$$

Since $R_\infty(i, j) = v(j)/[\phi v(i)]$, this simplifies to (11) after defining $\delta = \zeta/\phi$. \square

Thus, in principle, the realized return on *any* sufficiently long-dated fixed-income asset can proxy for the return on the long bond. For very large T , the decomposition in Result 11, which is related to results of Hansen and Scheinkman (2009), allows us to interpret pricing as the product of a time discount factor, ϕ^T , the economy-wide kernel, $v(i)$, which captures risk considerations, and a term specific to the asset, K .

More precisely, by Result 11, the difference in returns on maturity- T_1 assets and maturity- T_2 assets is of order $O(\delta^{\min\{T_1, T_2\}})$, where δ can be taken arbitrarily close to ψ/ϕ , the ratio of the second-largest and largest absolute values of the eigenvalues of \mathbf{A} . If, say, returns on 20-year and 30-year bonds are sufficiently far apart, then we can conclude that ψ/ϕ is close to 1. We will refer, in this case, to *slow convergence at the long end*.

Our next result characterizes a topological feature of the Markov chain driving the economy. Specifically, we will define a quantitative measure of the extent to which the economy features *traps*, and show how our measure can be computed empirically.

In order to state our key definition in a streamlined way, it will be convenient to write, for an arbitrary set of states $S \subseteq \{1, \dots, m\}$,

$$\mathbb{P}(\text{in } S) = \sum_{i \in S} \pi(i) \quad \text{and} \quad \mathbb{P}(\text{outside } S) = \sum_{i \notin S} \pi(i).$$

These represent the fraction of time the economy spends inside or outside S , respectively. We will also be interested in the fraction of time the economy spends exiting and entering S :

$$\mathbb{P}(\text{exit } S) = \sum_{\substack{i \in S \\ j \notin S}} \pi(i)\pi(i, j) \quad \text{and} \quad \mathbb{P}(\text{enter } S) = \sum_{\substack{i \notin S \\ j \in S}} \pi(i)\pi(i, j). \quad (12)$$

Finally, we define the conditional probabilities

$$\mathbb{P}(\text{exit } S \mid \text{start in } S) = \frac{\mathbb{P}(\text{exit } S)}{\mathbb{P}(\text{in } S)} \quad \text{and} \quad \mathbb{P}(\text{enter } S \mid \text{start outside } S) = \frac{\mathbb{P}(\text{enter } S)}{\mathbb{P}(\text{outside } S)}.$$

The first of these represents the probability that the economy exits the set of states S next period, conditional on starting inside S this period.

A *trap* is a collection of states that is hard to exit, once entered, so that $\mathbb{P}(\text{exit } S)$ is small in some appropriate sense. But we will want to rule out two ways in which $\mathbb{P}(\text{exit } S)$ can be small for trivial reasons. First, the set S may very small, so that the economy rarely exits S simply because it is rarely *in* S . We will deal with this issue by requiring that the probability of exiting S is small, conditional on starting in S . Second, if the set S is very large—almost the entire set of states $\{1, \dots, m\}$, say—then it will rarely be exited, but for an uninteresting reason.⁵ We deal with this case by requiring that the probability of entering S is small, conditional on starting outside S .

Definition 1. *The economy has an ε -trap if there is a collection of states, $S \subseteq \{1, \dots, m\}$, such that*

$$\mathbb{P}(\text{exit } S \mid \text{in } S) \leq \varepsilon \quad \text{and} \quad \mathbb{P}(\text{enter } S \mid \text{outside } S) \leq \varepsilon. \quad (13)$$

⁵There is another relatively uninteresting way in which an economy can have an ε -trap: if there are two states and the probability of exiting either one is ε then they are both ε -traps. We will see a more interesting example below. (It is possible, however, that the possibility we have labelled relatively uninteresting is relevant in practice.)

The property of having an ε -trap is stronger the closer ε is to zero; ε can therefore be thought of as an index of the extent to which a given economy experiences traps. A 1-trap is not a trap at all: every set of states satisfies (13) with $\varepsilon = 1$. A 0-trap is the extreme case with a set of states that can neither be escaped nor entered from outside: we have ruled out this possibility, in which the state space is disconnected, with our assumption that \mathbf{A} is a primitive matrix. We will sometimes say loosely that the economy *has a trap* if it has an ε -trap for small ε . A routine calculation shows that if S is an ε -trap, then the expected amount of time needed to escape the trap is at least $1/\varepsilon$.

If the matrix \mathbf{A} is known, the optimal (that is, smallest possible) value of ε for a given economy can in principle be computed mechanically by computing entry and exit probabilities for all subsets of states. This exercise may be computationally infeasible if there are many states, however—with, say, 100 states there are $2^{100} > 10^{30}$ subsets—and if the matrix \mathbf{A} is not observed, then this approach is not feasible even in principle.

Our next result therefore shows how ε can be linked to the data without direct knowledge of \mathbf{A} . It exploits the Cheeger inequality for directed graphs, which was proved by Chung (2005) and which, we believe, is new to the finance and economics literature.

Result 12. *Let \mathbf{x} be a vector that takes the values $x(1), \dots, x(m)$ in states $1, \dots, m$, and define*

$$\sigma^2(\Delta x) \equiv \sum_{i,j} \pi(i)\pi(i,j) [x(i) - x(j)]^2 \quad \text{and} \quad \sigma^2(x) \equiv \sum_i \pi(i)x(i)^2 - \left(\sum_i \pi(i)x(i) \right)^2,$$

so that $\sigma(\Delta x)$ is the volatility of changes in x and $\sigma(x)$ is the volatility of the level of x . Then the economy has an ε -trap with

$$\varepsilon = \frac{\sigma(\Delta x)}{\sigma(x)}.$$

Proof. Note first that the definition (13) of an ε -trap is equivalent to the condition that $\mathbb{P}(\text{exit } S) / \min \{\mathbb{P}(\text{in } S), 1 - \mathbb{P}(\text{in } S)\} \leq \varepsilon$. (This follows because $\mathbb{P}(\text{exit } S) = \mathbb{P}(\text{enter } S)$, as can be seen from the definition (12) using the fact that $\sum_i \pi(i)\pi(i, j) = \pi(j)$. Intuitively, it is clear that in a stationary model the long-run proportions of time spent entering and exiting any collection of states S must equal one another.)

By the Cheeger inequality for directed graphs (Chung, 2005, Theorem 5.1),

$$\frac{\lambda}{2} \leq \inf_S \frac{\mathbb{P}(\text{exit } S)}{\min \{\mathbb{P}(\text{in } S), 1 - \mathbb{P}(\text{in } S)\}} \leq \sqrt{2\lambda}, \quad (14)$$

where λ is the second-smallest eigenvalue of the Laplacian \mathcal{L} defined as

$$\mathcal{L} = \mathbf{I} - \frac{1}{2} \left(\mathbf{D}_\pi^{1/2} \mathbf{\Pi} \mathbf{D}_\pi^{-1/2} + \mathbf{D}_\pi^{-1/2} \mathbf{\Pi}' \mathbf{D}_\pi^{1/2} \right);$$

here \mathbf{D}_π is a diagonal matrix with the entries of $\boldsymbol{\pi}$ along its diagonal. Our interest is in the right-hand inequality in (14). By Corollary 4.2 of Chung (2005), λ satisfies

$$\lambda = \inf_x \sup_c \frac{\sum_{i,j} \pi(i)\pi(i, j) (x(i) - x(j))^2}{2 \sum_j \pi(j) (x(j) - c)^2}.$$

The inner supremum is attained if we set $c = \sum_k \pi(k)x(k)$, so we can rewrite

$$\lambda = \inf_x \frac{\sigma^2(\Delta x)}{2\sigma^2(x)},$$

with $\sigma^2(\Delta x)$ and $\sigma^2(x)$ as defined above. In conjunction with (14), this implies that there is a collection of states S such that

$$\frac{\mathbb{P}(\text{exit } S)}{\min \{\mathbb{P}(\text{in } S), 1 - \mathbb{P}(\text{in } S)\}} \leq \frac{\sigma(\Delta x)}{\sigma(x)}$$

for *all* vectors \boldsymbol{x} . The result follows. \square

Result 12 is very flexible, and can be applied using *any* vector \boldsymbol{x} that takes val-

ues $x(1), \dots, x(m)$. As $\frac{\sigma(\Delta x)}{\sigma(x)} = \sqrt{2[1 - \text{corr}(x_{t+1}, x_t)]}$, the relevant quantity is the maximal possible autocorrelation over all random variables that are measurable with respect to the state. Thus, given a collection of historical time series of variables relevant for fixed income pricing, the optimal choice of x_t (among linear combinations of the time series) is the *maximal autocorrelation factor* that can be constructed from the time series, as introduced by Switzer and Green (1984). This is somewhat similar to the first principal component in conventional principal component analysis (PCA), but has certain attractive properties. Unlike PCA, it is invariant to rescalings of the input time series. Moreover, it exploits the time-series nature of the data in a central way, whereas PCA generates the same factor definitions if the time series is randomly re-ordered.

To pursue one direction in which empirical work might proceed, suppose that we are given a collection of N time series $\mathbf{Z}(t) = \{Z_j(t)\}_{j=1, \dots, N}$, observed at times $t = 1, \dots, T$, that the $N \times N$ covariance matrix of the N time series is Σ , and that the $N \times N$ covariance matrix of the *differenced* time series is Σ_Δ . Then to find the optimal (i.e. lowest) possible value of ε among linear combinations of the time series $\mathbf{x} = \mathbf{w}'\mathbf{Z}$, we must choose \mathbf{w} to solve⁶

$$\min_{\mathbf{w}} \frac{\text{var}(\mathbf{w}'(\mathbf{Z}_{t+1} - \mathbf{Z}_t))}{\text{var}(\mathbf{w}'\mathbf{Z}_t)} = \min_{\mathbf{w}} \frac{\mathbf{w}'\Sigma_\Delta\mathbf{w}}{\mathbf{w}'\Sigma\mathbf{w}} \quad \text{or equivalently} \quad \min_{\mathbf{u}} \frac{\mathbf{u}'\Sigma^{-1/2}\Sigma_\Delta\Sigma^{-1/2}\mathbf{u}}{\mathbf{u}'\mathbf{u}},$$

where $\mathbf{u} = \Sigma^{1/2}\mathbf{w}$. The latter problem is solved by setting \mathbf{u} equal to the eigenvector of $\Sigma^{-1/2}\Sigma_\Delta\Sigma^{-1/2}$ with smallest eigenvalue (as the matrix is positive definite, all of its eigenvalues are positive). Having done so, $\mathbf{w} = \Sigma^{-1/2}\mathbf{u}$.

For example, if $\mathbf{Z} = (y_{3\text{mo}}, y_1, y_5, y_{10}, y_{20})'$, where $y_{3\text{mo}}, y_1, y_5, y_{10}$ and y_{20} are the 3-month, 1-year, 5-year, 10-year, and 20-year log yields from the St. Louis Fed's FRED database (observed monthly from October 1993 to June 2017), then the optimal choice of weights is $\mathbf{w} = (0.499, 0.285, 0.300, -1.569, 1.486)'$, with monthly autocorrelation

⁶See Haugen, Rajaratnam and Switzer (2015) for further detail on the maximum autocorrelation factor approach.

0.997 and implied $\varepsilon = 0.0716$, for an expected time trapped of at least $1/0.0716 = 14.0$ months. By contrast, the corresponding univariate calculations imply an expected time trapped of at least 12.1, 11.7, 8.3, 7.3, or 7.6 months for the 3-month, \dots , 20-year log yields, respectively.

Our next result links back to the speed of convergence at the long end.

Result 13. *Define χ to be the second-largest of the real parts of the eigenvalues of \mathbf{A} . Then the economy has an ε -trap where*

$$\varepsilon = \sqrt{2 \left(1 - \frac{\chi}{\phi}\right)}.$$

Proof. By the Cheeger inequality (14), there is an ε -trap for any $\varepsilon \geq \sqrt{2\lambda}$, where λ is the second-smallest eigenvalue of the Laplacian \mathcal{L} defined in the proof of Result 12. The result follows because

$$1 - \chi/\phi = \min(1 - \operatorname{Re} \rho_i) \geq \lambda,$$

where ρ_i are the eigenvalues of $\mathbf{\Pi}$ and the minimization is over all the eigenvalues of $\mathbf{\Pi}$ *other than the largest* (which equals one). The equality follows directly from the fact that the eigenvalues of \mathbf{A} equal the eigenvalues of $\mathbf{\Pi}$ multiplied by ϕ (as, by the decomposition (2) of Result 1, \mathbf{A}/ϕ and $\mathbf{\Pi}$ are similar matrices). The inequality follows from Theorem 4.3 of Chung (2005). \square

To interpret Result 13, note that if we restrict to economies in which all eigenvalues are real and positive⁷ (as in the example shown in Figure 1), then the eigenvalue-gap measures ψ/ϕ and χ/ϕ defined in Results 11 and 13 coincide. Then, slow convergence

⁷It is an empirical matter whether eigenvalues of \mathbf{A} —or equivalently $\mathbf{\Pi}$ —are real and positive in practice, but from a theoretical perspective it would be guaranteed if, say, $\mathbf{\Pi}$ were symmetric and the economy persistent in the sense that the diagonal entries of $\mathbf{\Pi}$ are larger than 1/2. [Proof: The eigenvalues of $2\mathbf{\Pi} - \mathbf{I}$ are real (by symmetry) and lie in $(-1, 1]$ (by the Perron–Frobenius theorem, $2\mathbf{\Pi} - \mathbf{I}$ being stochastic). It follows that the eigenvalues of $\mathbf{\Pi}$ lie in $(0, 1]$.]

at the long end implies the existence of a trap. Even without the assumption that eigenvalues are real and positive, we can still say that if $\chi/\phi \approx 1$ —so that there is a trap with small ε , by Result 13—then $\psi/\phi \approx 1$, because $\psi \geq \chi$; so that convergence at the long end is slow, by Result 11.

An example.—Consider two economies. There are eight states of the world in each, and the short rate is 0% in states 1 through 4, and 10% in states 5 through 8.⁸ But the evolution of the state variable differs across the two economies.




Figure 2: The state space in Economy A (left) and Economy B (right). The short rate is 0% in states 1 through 4 and 10% in states 5 through 8. There is a 50% probability of a transition along any arrow. The excess return on the long bond in each possible transition (in %) is indicated at the mid point of each arrow.

In *Economy A*, the economy transitions with equal probability between any two states. Thus states 1 through 4 are essentially indistinguishable, and can be compressed into one super-state, as shown in Figure 2; and similarly for states 5 through 8. In *Economy B*, the two super-states have a non-trivial internal structure. Starting in state 1, the economy remains in state 1 or transitions to state 2. From state 2, the economy transitions to state 3, or back to state 1; from state 3, to state 4 or back to state 1; from state 4, to state 5—in which case the short rate jumps to 10%—or back to state 1 (yet again). The situation is symmetrical in states 5 through 8 (with state 6 playing the corresponding role to state 1). Thus interest rates get stuck for

⁸The Arrow–Debreu matrices \mathbf{A} that represent asset prices in the two economies are given in the Appendix.

extended periods. States 1 through 4 can be thought of as a liquidity trap, while states 5 through 8 represent a high-interest-rate regime.




Figure 3: Yield curves in Economy A (left) and Economy B (right).

The yield curves in each economy are shown in Figure 3. In Economy A, yields converge fairly rapidly to the long rate. In Economy B, the long end of the yield curve does not converge (in the sense of approximating the infinitely long yield) at horizons that are plausibly observable; and in the depths of a liquidity trap—in state 1—the yield curve is flat, at zero, over a range of shorter maturities. Yields are more variable across states and over time. The risk metric, Q , of Result 10 is therefore substantially larger than in Economy A. Correspondingly, the long bond’s excess return (indicated at the mid point of each of the arrows in Figure 2) is substantially more volatile in Economy B than in Economy A. The long bond’s Sharpe ratio is constant, and just below 5%, in Economy A, whereas it is more than 15% on average in Economy B and state-dependent (with a conditional Sharpe ratio of 39%, 60%, and 41% in states 3, 4, and 5 respectively).

If we let S be the collection of states 1 through 4, then⁹

$$\mathbb{P}^A(\text{exit } S \mid \text{start in } S) = \frac{1}{2} \quad \text{and} \quad \mathbb{P}^A(\text{enter } S \mid \text{start outside } S) = \frac{1}{2}$$

⁹This choice of S is optimal for each economy, as is easily checked. The stationary distributions in each economy are given by the positive left eigenvectors of the transition matrices, namely $\pi'_A = (\frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8})'$ and $\pi'_B = (\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{16}, \frac{1}{16}, \frac{1}{4}, \frac{1}{8}, \frac{1}{8})'$.

and

$$\mathbb{P}^B(\text{exit } S \mid \text{start in } S) = \frac{1}{16} \quad \text{and} \quad \mathbb{P}^B(\text{enter } S \mid \text{start outside } S) = \frac{1}{16},$$

so there is an ε -trap with $\varepsilon = 1/2$ in Economy A, and with $\varepsilon = 1/16$ in economy B. The latter, lower, value formalizes the sense in which economy B is more trapped than economy A.

Finally, we note that Economy A is *time-reversible*: it looks the same whether run forwards or backwards in time. In contrast, there is an arrow of time in Economy B (for example, the economy never transitions from state 3 to state 2). This is a feature of our framework that is not shared by stationary Gaussian models, which are time-reversible (Weiss, 1975).

4 What is needed for recovery?

Most of our results thus far have avoided the need to observe \mathbf{A} directly: we have shown, for example, how to infer the time preference rate ϕ , the kernel \mathbf{v} , and the excess return on the long bond, from specific—if idealized—asset prices. To recover the probability matrix $\mathbf{\Pi}$, however, we must assume (as Ross, 2015, did) that the entire matrix of Arrow–Debreu prices, \mathbf{A} , is directly observable. But is it possible to infer \mathbf{A} from asset prices that are clearly easily observable—for example, from yield curve information alone?

A related concern is that it is not clear that the econometrician observes asset prices in all states of the world. Can the cross-section substitute for the time series? To sharpen the question: can \mathbf{A} can be inferred from perfect knowledge of all Arrow–Debreu prices, at all maturities, in a single state?¹⁰

We now address these questions by considering some intentionally stylized exam-

¹⁰This possibility was outlined by Ross (2015) and explored more formally by Jensen, Lando and Pedersen (2018).

ples. These examples are as simple as we could make them, with few states of the world and Arrow–Debreu prices set to zero wherever possible. But they illustrate the flexibility of our framework and exhibit some potential pitfalls for empirical work.

We will repeatedly exploit the following fact. Suppose we observe the prices, in each state, of assets with payoff vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$: that is, we know $\mathbf{A}\mathbf{x}_1, \dots, \mathbf{A}\mathbf{x}_m$. Then this information uniquely determines \mathbf{A} if and only if $\mathbf{x}_1, \dots, \mathbf{x}_m$ are linearly independent (i.e., span \mathbb{R}^m). As a trivial example, if \mathbf{e}_j denotes the j th unit vector then knowledge of $\mathbf{A}\mathbf{e}_1$ (that is, of the price, in each state, of the Arrow–Debreu security that pays off in state 1), of $\mathbf{A}\mathbf{e}_2$, and so on, is sufficient to reveal \mathbf{A} .

For a less trivial example, suppose that yields out to maturity m are observable in every state. Then we need the vectors $\mathbf{x}_1 = \mathbf{e}$, $\mathbf{x}_2 = \mathbf{A}\mathbf{e}$, \dots , $\mathbf{x}_m = \mathbf{A}^{m-1}\mathbf{e}$ —that is, bond prices at maturities from 0 to $m - 1$ —to be linearly independent.

The question is whether this assumption is plausible. In practice, it might be the case that (say) a 4-year zero-coupon bond trades close to the midpoint of 3- and 5-year zero-coupon bond prices in every state of the world; if so, $\mathbf{A}^3\mathbf{e}$, $\mathbf{A}^4\mathbf{e}$, and $\mathbf{A}^5\mathbf{e}$ will be approximately collinear even if not perfectly linearly dependent. More generally, the literature has documented that bond yields approximately obey a low-dimensional factor structure. In such circumstances, estimates of \mathbf{A} may be unstable in practice.

Result 14 (What can go wrong). *Suppose we observe bond prices at maturities 1, \dots , m in every state. If $\mathbf{e}, \dots, \mathbf{A}^{m-1}\mathbf{e}$ (that is, the vectors of bond prices at maturities up to $m - 1$, together with the vector of ones) are linearly independent then \mathbf{A} is identified. Otherwise it is not—even if we observe bond prices at all maturities in every state.*

Proof. If the given collection of vectors is linearly independent then it is a basis for \mathbb{R}^m , and any matrix can be identified from its action on a basis. Conversely, suppose $\mathbf{e}, \dots, \mathbf{A}^{m-1}\mathbf{e}$ are linearly dependent. Then $\mathbf{e}, \dots, \mathbf{A}^N\mathbf{e}$ are linearly dependent for any $N \geq m - 1$ (and hence longer maturities will not help, as we do not observe the action of \mathbf{A} on any basis). This follows because, by the Cayley–Hamilton theorem,

there is a linear dependence between the matrix \mathbf{A}^m and the matrices \mathbf{I} , \mathbf{A} , \dots , \mathbf{A}^{m-1} . \square

Example 1 (Nonrecoverability despite perfect knowledge of yield curves in all states).— Knowledge of the entire yield curve in every state of the world—that is, knowledge of bond prices at all maturities in all states—need not determine \mathbf{A} . The simplest possible nontrivial illustration of this fact requires at least three states (for with $m = 2$ states, nonrecoverability of \mathbf{A} from knowledge of yield curves requires that \mathbf{e} and $\mathbf{A}\mathbf{e}$ are linearly dependent, and hence that interest rates are constant). Thus, consider two economies with different Arrow–Debreu price matrices,

$$\mathbf{A}_1 = \begin{pmatrix} 0 & 0.5 & 0.5 \\ 0.45 & 0 & 0.5 \\ 0.4 & 0.5 & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{A}_2 = \begin{pmatrix} 0.25 & 0 & 0.75 \\ 0 & 0.9 & 0.05 \\ 0.4 & 0.5 & 0 \end{pmatrix}.$$

The row sums of each matrix are 1, 0.95 and 0.9, so in each economy the one-period interest rate is zero in state 1, about 5% in state 2, and about 10% in state 3; thus knowledge of the values taken by the short rate does not distinguish between the two economies. In fact an easy calculation shows that the *entire yield curve* is identical for both economies in every state: see Figure 4a. And yet the risk premium on the long bond differs across the two economies in every state. (We compute the bond risk premium using the probabilities recovered in Proposition 1.) The key property of the example is that the vectors \mathbf{e} , $\mathbf{A}_i\mathbf{e}$ and $\mathbf{A}_i^2\mathbf{e}$ are not linearly independent for $i = 1$ or 2 , so do not constitute a basis for \mathbb{R}^3 .

Note, however, that by Result 1, if the yield curve is observed in every state then we can at least observe ϕ and \mathbf{v} from the evolution of the long bond. What we *cannot* infer, if \mathbf{A} itself is not known, is the transition matrix $\mathbf{\Pi}$.

If bond yields alone are not enough, will a richer cross-section of asset prices permit identification of \mathbf{A} ? The answer, trivially, is yes if the prices of all one-period

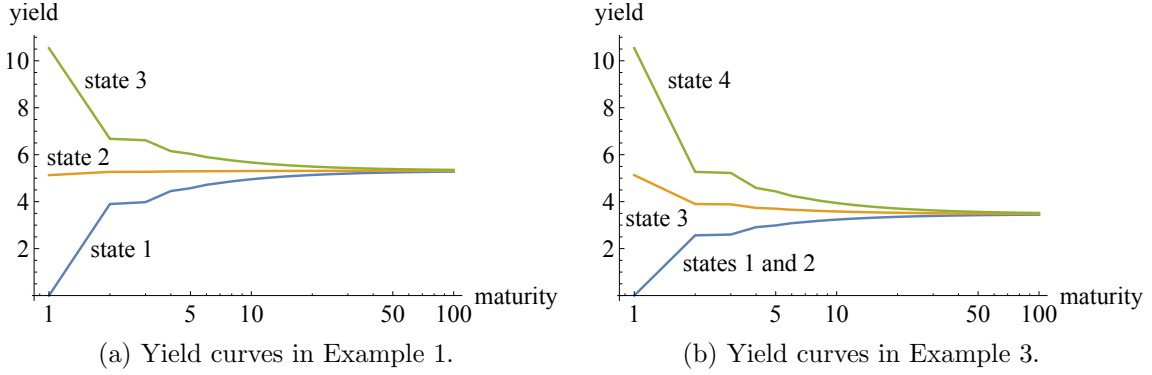


Figure 4: In Example 1, two different economies generate the same yield curves in each of the states, but risk premia differ in the two economies. In Example 3, there is just one economy; the yield curve is identical in states 1 and 2, but bond risk premia are different in the two states. (The jaggedness of the yield curves is inessential: it arises because we chose to give examples that are easily described in words.)

Arrow–Debreu securities are observed in all states. But it is not necessarily enough to observe all Arrow–Debreu securities in a *single* state, even if they are observed at all possible maturities.

Result 15 (What can go wrong, part 2). *Suppose we observe the prices of all Arrow–Debreu securities at maturities $1, \dots, m$ in a single state—call it state 1. This is equivalent to observing $\mathbf{A}'\mathbf{e}_1, \dots, (\mathbf{A}')^m\mathbf{e}_1$. If $\mathbf{e}_1, \dots, (\mathbf{A}')^{m-1}\mathbf{e}_1$ span \mathbb{R}^m then \mathbf{A} is identified. Otherwise it is not, even if we observe (in state 1) the prices of all Arrow–Debreu securities at all maturities.*

Proof. As in the preceding result, if $\mathbf{e}_1, \dots, (\mathbf{A}')^{m-1}\mathbf{e}_1$ span \mathbb{R}^m then they are a basis and we observe the action of \mathbf{A}' on this basis. Thus \mathbf{A}' is identified, and hence also \mathbf{A} . The converse direction also proceeds as before: by the Cayley–Hamilton theorem, if $\mathbf{e}_1, \dots, (\mathbf{A}')^{m-1}\mathbf{e}_1$ are linearly dependent then so are $\mathbf{e}_1, \dots, (\mathbf{A}')^N\mathbf{e}_1$ for any N . \square

Example 2 (Nonrecoverability from all asset prices in a single state).—Suppose we observe, in state 1, the prices of every possible Arrow–Debreu security at every possible maturity. Result 15 shows that \mathbf{A} is identified if and only if $\mathbf{e}_1, \dots, (\mathbf{A}')^{m-1}\mathbf{e}_1$

are linearly independent. For a simple example in which this fails, let

$$\mathbf{A}_3 = \begin{pmatrix} 0 & 0.5 & 0.5 \\ 0.45 & 0.5 & 0 \\ 0.4 & 0 & 0.5 \end{pmatrix}.$$

In state 1, the price of a state- i Arrow–Debreu security maturing at time T is the same in this economy as in the economy \mathbf{A}_1 of Example 1 for every i and T . (As a corollary, the yield curve is identical in state 1 in the two economies.) Hence, even given perfect knowledge of all possible asset prices in state 1, the matrix \mathbf{A} cannot be identified.

The proofs of Results 14 and 15 exploited the Cayley–Hamilton theorem, which implies that \mathbf{A}^m can be written as a linear combination of lower order powers of \mathbf{A} .¹¹ But there may be an even lower order relationship between powers of \mathbf{A} . In this case \mathbf{A} is neither identifiable from the term structure (as in Result 14) nor from all Arrow–Debreu prices in a single state (as in Result 15).

Example 3 (What can go wrong: the quadratic case).—The simplest nontrivial example¹² of this phenomenon—and, in a sense, the worst-case scenario, because the dependence is of the lowest possible order—arises if $\mathbf{A}^2 = \phi\mathbf{A}$ for some ϕ . Then $(\mathbf{A}/\phi)^2 = \mathbf{A}/\phi$, so \mathbf{A}/ϕ is a projection matrix. This projection must be onto a *line*: a projection onto a higher-dimensional subspace would have a repeated maximal eigenvalue equal to 1, which is not possible by the Perron–Frobenius theorem (\mathbf{A} being primitive). Such a projection can be written in the form $\mathbf{A}/\phi = \mathbf{v}\mathbf{v}'$, where \mathbf{v} is a unit vector (i.e., satisfies $\mathbf{v}'\mathbf{v} = 1$) that determines the line onto which the

¹¹As matrices of dimension m inhabit an m^2 -dimensional vector space, the matrices \mathbf{I} , \mathbf{A} , \dots , \mathbf{A}^{m^2} are linearly dependent. The Cayley–Hamilton theorem strengthens this fact considerably.

¹²The two simpler cases—linear dependence between \mathbf{A} and \mathbf{I} , or between \mathbf{A}^2 and \mathbf{I} —are inconsistent with primitivity of \mathbf{A} . For example, $\mathbf{A}^2 = \lambda\mathbf{I}$ implies that all even powers of \mathbf{A} have zero off-diagonal terms. More generally, note that matrix quadratic equations exhibit richer properties than their scalar counterparts. For example the equation $\mathbf{A}^2 = \mathbf{I}$ has uncountably many solutions even in the case $m = 2$.

projection takes place. The converse also holds: if the Arrow–Debreu price matrix satisfies

$$\mathbf{A} = \phi \mathbf{v} \mathbf{v}', \quad (15)$$

where $\mathbf{v}'\mathbf{v} = 1$ and we require that $\phi > 0$ and \mathbf{v} has all positive entries to avoid arbitrage opportunities, then $\mathbf{A}^2 = \phi \mathbf{A}$ and (consistent with previous notation) \mathbf{v} is an eigenvector of \mathbf{A} with eigenvalue ϕ .

The problem of recovery is then particularly stark: as $\mathbf{A}^t = \phi^{t-1} \mathbf{A}$ for all t , there is no information in the term structure of *any* asset that is not present in the price of a short-dated asset (other than learning the value of ϕ itself which, as always, can be inferred from the long yield in any state). Moreover, the price of a t -period bond is $\mathbf{A}^t \mathbf{e} = \phi^{t-1} \mathbf{A} \mathbf{e}$ so in every state of the world, i , the yield curve satisfies

$$y_t(i) - y_\infty = \frac{1}{t}(y_1(i) - y_\infty).$$

The characterization (15) and Result 1 together imply that $\mathbf{D} \Pi \mathbf{D}^{-1} = \mathbf{v} \mathbf{v}'$, where \mathbf{D} has the vector \mathbf{v} along its diagonal. This implies that $\pi_{i,j} = v_j^2$, so that the probability of transitioning to state j is the same in every state: that is, the economy evolves in an i.i.d. fashion.

Now suppose that \mathbf{A} obeys a more general quadratic equation. Its eigenvalues must satisfy the scalar version of the same equation so there are two, ϕ_1 and ϕ_2 , and by the Perron–Frobenius theorem they must satisfy $\phi_1 > |\phi_2|$, with the former having multiplicity 1 and the latter multiplicity $m - 1$. As ϕ_1 is real, ϕ_2 must also be real, because the eigenvalue sum equals the trace of \mathbf{A} (which is manifestly real).

We can therefore write $\mathbf{A}^2 - (\phi_1 + \phi_2)\mathbf{A} + \phi_1\phi_2\mathbf{I} = \mathbf{0}$, or equivalently

$$\left(\frac{\mathbf{A} - \phi_2 \mathbf{I}}{\phi_1 - \phi_2} \right)^2 = \frac{\mathbf{A} - \phi_2 \mathbf{I}}{\phi_1 - \phi_2}.$$

It follows that $\frac{\mathbf{A} - \phi_2 \mathbf{I}}{\phi_1 - \phi_2}$ is a projection matrix with 0 and 1 as its only eigenvalues.

Moreover, it can only have one maximal eigenvalue (because \mathbf{A} has one maximal eigenvalue) so represents a projection onto a *line*. Thus $\frac{\mathbf{A}-\phi_2\mathbf{I}}{\phi_1-\phi_2} = \mathbf{v}\mathbf{v}'$, that is, $\mathbf{A} = (\phi_1 - \phi_2)\mathbf{v}\mathbf{v}' + \phi_2\mathbf{I}$, where $\mathbf{v}'\mathbf{v} = 1$ (and \mathbf{v} is the Perron eigenvector).

In conjunction with Result 1 this implies that

$$\frac{\phi_1 - \phi_2}{\phi_1}\mathbf{v}\mathbf{v}' = \mathbf{D} \left(\mathbf{\Pi} - \frac{\phi_2}{\phi_1}\mathbf{I} \right) \mathbf{D}^{-1}.$$

As \mathbf{D} is diagonal (with \mathbf{v} along the diagonal) it follows that $\pi_{i,j} = \frac{\phi_1-\phi_2}{\phi_1}v_j^2 + \frac{\phi_2}{\phi_1}\mathbf{1}_{i=j}$. This is similar to the i.i.d. case, as in the first part of the present example, but with transition probabilities distorted to allow for stickiness in (if $\phi_2 > 0$) or repulsion from (if $\phi_2 < 0$) the current state.

Our final example makes a different point. The yield curve distinguishes between states in Examples 1 and 2, so can act as the state variable. If, in addition to price information, one also observes time series information about the evolution of the yield curve, then the transition matrix $\mathbf{\Pi}$ can be determined in principle, and hence identification of \mathbf{A} is possible (in fact, by Result 1, we would only need to observe the time-series properties of the long end of the yield curve and hence ϕ and \mathbf{v}). But the yield curve is not a suitable state variable in general.

Example 4 (Hidden factors; and the importance of lags).—Two different states may have different bond risk premia but identical yield curves. Suppose that

$$\mathbf{A} = \begin{pmatrix} 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0.45 & 0 & 0.5 & 0 \\ 0 & 0.9 & 0 & 0 \end{pmatrix}.$$

Figure 4b plots the yield curve in each of the four states. The short rate is 0% in states 1 and 2; about 5% in state 3; and about 10% in state 4.

We have chosen to specify that \mathbf{A} has several zero entries so that it is possible to

give a simple description of the states:

- In state 1, interest rates are low and volatility is high. The economy may either remain in state 1 or transition to state 4. The risk premium on the long bond is positive.
- In state 2, interest rates are low, and the economy transitions deterministically to state 3. The risk premium on the long bond is zero.
- In state 3, interest rates are intermediate and volatility is high. The economy may either remain in state 3 or transition to state 1. The risk premium on the long bond is positive.
- In state 4, interest rates are high, and the economy transitions deterministically to state 2. The risk premium on the long bond is zero.

(The presence of the zero entries in \mathbf{A} explains the jagged shapes of the yield curves in Figure 4b. For an example that generates smooth yield curves, raise \mathbf{A} to the fifth¹³ power—and, if desired, multiply by a scalar—to generate a matrix whose entries are all positive and which has all the relevant properties exhibited in this example.)

As in Example 1, \mathbf{e} , $\mathbf{A}\mathbf{e}$, $\mathbf{A}^2\mathbf{e}$, and $\mathbf{A}^3\mathbf{e}$ are linearly dependent, so knowledge of the yield curve in every state is not sufficient to recover \mathbf{A} . In fact, we have set up the example so that the yield curve is identical in states 1 and 2. But the long bond risk premium is positive in state 1 and zero in state 2: thus bond prices need not “span” the uncertainty in fixed income markets. Duffee (2011) and Joslin, Priebisch and Singleton (2014) have argued (in a Gaussian context) that this is the empirically relevant case.

The example—though stylized—also illustrates another interesting phenomenon. While the yield curve itself is not a satisfactory state variable, knowledge of the current *and lagged* yield curves reveals the state perfectly (the economy is in state 2 if

¹³Our maintained assumption that \mathbf{A} is primitive ensures that above some time horizon T all Arrow–Debreu prices are strictly positive; in this example $T = 5$ is enough.

and only if it was in state 4 last period). Intriguingly, Cochrane and Piazzesi (2005) find that information in the lagged term structure is indeed useful in forecasting bond risk premia. They interpret their empirical finding as evidence of measurement error, and write, “Bond prices are time- t expected values of future discount factors, so a full set of time- t bond yields should drive out lagged yields in forecasting regressions. . . Bond prices *reveal* all other important state variables. For this reason, term structure models do not include lags.” Example 4 shows that lags may in fact assist in forecasting even if yields are perfectly observed.¹⁴

5 Conclusions

We have studied the framework of Ross (2015), in which the state of the economy follows a discrete-time, finite-state Markov chain and markets are complete; and shown that recovery can be partially effected by studying the long end of the yield curve, without knowledge of the full matrix of Arrow–Debreu prices which Ross assumed to be directly observable. More precisely, we show that it is possible to infer the time discount factor and marginal utilities of what we call a pseudo-representative investor from the behavior of the long bond alone.

Our results place Ross’s Recovery Theorem in a broader context that has been explored by authors including Backus, Gregory and Zin (1989), Kazemi (1992), Bansal and Lehmann (1994, 1997), Alvarez and Jermann (2005) and Hansen and Scheinkman (2009), and clarify that the key property implied by Ross’s structural assumptions is that the long bond is growth-optimal relative to the set of assets under consideration.

We acknowledge that if the goal is recovery in equity markets, this property is implausible (as argued by Borovička, Hansen and Scheinkman, 2017). We suggest, however, that the critique has less force in the context of fixed income markets, where the relevant state variables are more plausibly stationary. If one restricts

¹⁴Expanding the set of observables in other ways can also help. For example, states 1 and 2 could also be distinguished from one another using the implied volatility of short-dated bonds.

attention to fixed income markets, it remains an open empirical question whether the set of probabilities that emerge in the decomposition of Result 1 are in fact the true probabilities—or, equivalently, whether the long bond is growth-optimal among fixed-income assets. Assuming that the recovered probabilities *are* the true probabilities, various interesting facts follow, notably that the yield curve is upward-sloping on average, and that long bond option prices reveal expected returns on the long bond.

There are nontrivial empirical issues that must be confronted if our theoretical results are to be implemented in practice. Most notably, the speed of convergence of long-but-finite bonds to the idealized long bond is of central importance. We derive various results that bear on this issue, and which we hope are interesting in their own right: we introduce, for example, a measure of the “trappedness” of an economy and relate it to the speed of convergence at the long end of the yield curve and to the eigenvalue gap (between the largest and second-largest eigenvalues) of \mathbf{A} .

Aside from issues related to recovery, our framework is well suited to studying situations in which interest rates exhibit cycles or traps; or, more generally, to cases in which the topology of the state space is nontrivial (whether due to reputational considerations on the part of monetary policymakers, to liquidity traps, to technological or other forms of irreversibility, or to something else). As a result it can be used to address certain issues that are assumed away in Gaussian models. Stationary Gaussian models are time-reversible (Weiss, 1975), for example, so the conclusions reached by an econometrician living in a stationary Gaussian world would be the same whether time runs forwards or backwards. Such models exclude the possibility that, say, interest rates “go up by the stairs and down by the elevator” (or the converse). This is a highly restrictive assumption, and one that is empirically dubious.¹⁵

¹⁵See, for example, Neftçi (1984) and Ramsey and Rothman (1996).

6 References

- Alvarez, F., and U. J. Jermann (2005), “Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth,” *Econometrica*, 73:6:1977–2016.
- Backus, D. K., A. W. Gregory, and S. E. Zin (1989), “Risk Premiums in the Term Structure: Evidence from Artificial Economies,” *Journal of Monetary Economics*, 24:371–399.
- Bakshi, G., F. Chabi-Yo, and X. Gao (2018), “A Recovery that We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem,” *Review of Financial Studies*, 31:2:532–555.
- Bansal, R., and B. N. Lehmann (1994), Addendum to “Growth-Optimal Portfolio Restrictions on Asset Pricing Models,” working paper.
- Bansal, R., and B. N. Lehmann (1997), “Growth-Optimal Portfolio Restrictions on Asset Pricing Models,” *Macroeconomic Dynamics*, 1:333–354.
- Borovička, J., L. Hansen and J. Scheinkman (2017), “Misspecified Recovery,” *Journal of Finance*, 71:6:2493–2544.
- Breeden, D. T., and R. H. Litzenberger (1978), “Prices of State-Contingent Claims Implicit in Option Prices,” *Journal of Business*, 51:4:621–651.
- Chung, F. (2005), “Laplacians and the Cheeger Inequality for Directed Graphs,” *Annals of Combinatorics*, 9:1–19.
- Cochrane, J. H., and M. Piazzesi (2005), “Bond Risk Premia,” *American Economic Review*, 95:1:138–160.
- Duffee, G. R., (2011), “Information in (and not in) the Term Structure,” *Review of Financial Studies*, 24:9:2895–2934.
- Dybvig, P. H., J. E. Ingersoll, and S. A. Ross (1996), “Long Forward and Zero-Coupon Rates Can Never Fall,” *Journal of Business*, 69:1:1–25.
- Geman, H., N. El Karoui and J.-C. Rochet (1995), “Changes of Numéraire, Changes of Probability Measure and Option Pricing,” *Journal of Applied Probability*, 32:2:443–458.

- Hansen, L. P., and J. A. Scheinkman (2009), “Long-Term Risk: An Operator Approach,” *Econometrica*, 77:1:177–234.
- Haugen, M. A., B. Rajaratnam, and P. Switzer (2015), “Extracting Common Time Trends from Concurrent Time Series: Maximum Autocorrelation Factors with Applications,” working paper.
- Horn, R. A., and C. R. Johnson (1990), *Matrix Analysis*, Cambridge University Press, Cambridge, UK.
- Jensen, C. S., D. Lando and L. H. Pedersen (2018), “Generalized Recovery,” *Journal of Financial Economics*, forthcoming.
- Joslin, S., M. Priebsch, and K. J. Singleton (2014), “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” *Journal of Finance*, 69:3:1197–1233.
- Kazemi, H. B. (1992), “An Intemporal Model of Asset Prices in a Markov Economy with a Limiting Stationary Distribution,” *Review of Financial Studies*, 5:1:85–104.
- Kremens, L., and I. W. R. Martin (2018), “The Quanto Theory of Exchange Rates,” *American Economic Review*, forthcoming.
- Martin, I. W. R. (2017), “What is the Expected Return on the Market?,” *Quarterly Journal of Economics*, 132:1:367–322.
- Martin, I. W. R., and C. Wagner (2018), “What is the Expected Return on a Stock?,” *Journal of Finance*, forthcoming.
- Neftçi, S. N. (1984), “Are Economic Time Series Asymmetric over the Business Cycle?,” *Journal of Political Economy*, 92:2:307–328.
- Ramsey, J. B., and P. Rothman (1996), “Time Irreversibility and Business Cycle Asymmetry,” *Journal of Money, Credit and Banking*, 28:1:1–21.
- Ross, S. A. (2015), “The Recovery Theorem,” *Journal of Finance*, 70:2:615–648.
- Switzer, P., and A. A. Green (1984), “Min/Max Autocorrelation Factors for Multivariate Spatial Imagery,” Technical Report No. 6, Department of Statistics, Stanford University.
- Weiss, G. (1975), “Time-Reversibility of Linear Stochastic Processes,” *Journal of Ap-*

plied Probability, 12:4:831–836.

A Appendix

The Arrow–Debreu price matrices for the example at the end of Section 3 are

$$\begin{pmatrix} 0.119 & 0.119 & 0.119 & 0.119 & 0.131 & 0.131 & 0.131 & 0.131 \\ 0.119 & 0.119 & 0.119 & 0.119 & 0.131 & 0.131 & 0.131 & 0.131 \\ 0.119 & 0.119 & 0.119 & 0.119 & 0.131 & 0.131 & 0.131 & 0.131 \\ 0.119 & 0.119 & 0.119 & 0.119 & 0.131 & 0.131 & 0.131 & 0.131 \\ 0.107 & 0.107 & 0.107 & 0.107 & 0.119 & 0.119 & 0.119 & 0.119 \\ 0.107 & 0.107 & 0.107 & 0.107 & 0.119 & 0.119 & 0.119 & 0.119 \\ 0.107 & 0.107 & 0.107 & 0.107 & 0.119 & 0.119 & 0.119 & 0.119 \\ 0.107 & 0.107 & 0.107 & 0.107 & 0.119 & 0.119 & 0.119 & 0.119 \end{pmatrix}$$

in Economy A, and

$$\begin{pmatrix} 0.462 & 0.538 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.396 & 0 & 0.604 & 0 & 0 & 0 & 0 & 0 \\ 0.303 & 0 & 0 & 0.697 & 0 & 0 & 0 & 0 \\ 0.201 & 0 & 0 & 0 & 0.799 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.267 & 0 & 0.638 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.462 & 0.443 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.481 & 0 & 0.423 \\ 0 & 0 & 0 & 0 & 0.380 & 0.525 & 0 & 0 \end{pmatrix}$$

in Economy B.