

Generalized Recovery

Jensen, Christian Skov ; Lando, David; Pedersen, Lasse Heje

Document Version
Final published version

Published in:
Journal of Financial Economics

DOI:
[10.1016/j.jfineco.2018.12.003](https://doi.org/10.1016/j.jfineco.2018.12.003)

Publication date:
2019

License
CC BY-NC-ND

Citation for published version (APA):
Jensen, C. S., Lando, D., & Pedersen, L. H. (2019). Generalized Recovery. *Journal of Financial Economics*, 133(1), 154-174. <https://doi.org/10.1016/j.jfineco.2018.12.003>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 22. Mar. 2025





Generalized recovery

Christian Skov Jensen^{a,*}, David Lando^{b,e}, Lasse Heje Pedersen^{b,c,d,e}

^a Bocconi University, Via Roentgen 1, Milan 20136, Italy

^b Copenhagen Business School, Solbjerg Plads 3, Frederiksberg 2000, Denmark

^c Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, NY 10012, USA

^d AQR Capital Management, CT 06830, USA

^e Center for Economic Policy Research (CEPR), London, UK

ARTICLE INFO

Article history:

Received 20 February 2017

Revised 22 June 2017

Accepted 21 July 2017

Available online 22 December 2018

JEL classification:

G1

G12

G13

G17

Keywords:

Recovery

Asset pricing

Pricing kernel

Predicting returns

ABSTRACT

We characterize when physical probabilities, marginal utilities, and the discount rate can be recovered from observed state prices for several future time periods. We make no assumptions of the probability distribution, thus generalizing the time-homogeneous stationary model of Ross (2015). Recovery is feasible when the number of maturities with observable prices is higher than the number of states of the economy (or the number of parameters characterizing the pricing kernel). When recovery is feasible, our model allows a closed-form linearized solution. We implement our model empirically, testing the predictive power of the recovered expected return and other recovered statistics.

© 2019 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license.
(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

The holy grail in financial economics is to decode probabilities and risk preferences from asset prices. This decoding has been viewed as impossible until Ross (2015) provided sufficient conditions for such a recovery in a time-homogeneous Markov economy (using the Perron–Frobenius theorem). However, his recovery method has been criticized by Borovicka et al. (2016) (who also rely on Perron–Frobenius and results of Hansen and Scheinkman 2009), arguing that Ross's assumptions rule out realistic models.

This paper sheds new light on this debate, both theoretically and empirically. Theoretically, we generalize the


recovery theorem to handle a general probability distribution that makes no assumptions of time-homogeneity or Markovian behavior. We show when recovery is possible—and when it is not—using a simple counting argument (formalized based on Sard's theorem), which focuses the attention on the economics of the problem. When recovery is possible, we show that our recovery inversion from prices to probabilities and preferences can be implemented in closed form. We implement our method empirically using option data from 1996 to 2015 and study how the recovered expected returns predict future actual returns.

To understand our method, note first that Ross (2015) assumes that state prices are known not just in each final state but also starting from each possible current state as illustrated in Fig. 1, Panel A. Simply put, he assumes that we know all prices today and all prices in all “parallel universes” with different starting points. Since we clearly cannot observe such parallel universes, Ross (2015)


* Corresponding author.

E-mail address: christian.jensen@unibocconi.it (C.S. Jensen).

Panel A. Ross’s recovery theorem: one period, two “parallel universes”



Panel B. Ross’s recovery theorem: time-homogeneous dynamic setting



Panel C. Our generalized recovery: no assumptions about probabilities




Fig. 1. Generalized recovery framework. Panel A illustrates the idea behind Ross’s recovery theorem, namely that we start with information about all Arrow–Debreu prices in *all* initial states (not just the state we are currently in, but also prices in “parallel universes” where today’s state is different). Panel B shows how Ross moves to a dynamic setting by assuming time-homogeneity, that is, assuming that the prices and probabilities are the same for the two dotted lines, and so on for each of the other pairs of lines. Panel C illustrates our generalized recovery method, where we make no assumptions about the probabilities.

proposes to implement his model based on prices for several future time periods, relying on the assumption that all time periods have identical structures for prices and probabilities (time-homogeneity), illustrated in Fig. 1, Panel B. In other words, Ross assumes that if S&P 500 is at the level 2000, then one-period option prices do not depend on the calendar time at which this level is observed.

We show that the recovery problem can be simplified by starting directly with the state prices for all future times given only the *current* state (Fig. 1, Panel C). We impose no dynamic structure on the probabilities, allowing the probability distribution to be fully general at each future time, thus relaxing Ross’s time-homogeneity assumption that is unlikely to be met empirically.

We first show that when the number of states S is no greater than the number of time periods T , then recovery is possible. To see the intuition, consider simply

the number of equations and the number of unknowns. First, we have S equations at each time period, one for each Arrow–Debreu price, for a total of ST equations. Second, we have 1 unknown discount rate, $S - 1$ unknown marginal utilities, and $S - 1$ unknown probabilities for each future time period. In conclusion, we have ST equations with $1 + (S - 1) + (S - 1)T = ST + S - T$ unknowns. These equations are not linear, but we provide a precise sense in which we can essentially just count equations. Hence, recovery is possible when $S \leq T$.

To understand the intuition behind this result, note that, for each time period, we have S equations and only $S - 1$ probabilities. Hence, for each additional time period we have one extra equation that can help us recover the marginal utilities and discount rate, and the number of marginal utilities does not grow with the number of time periods.

By focusing on square matrices, Ross's model falls into the category $S = T$, so our counting argument explains why he finds recovery. However, our method applies under much more general conditions. We show that when Ross's time-homogeneity conditions are met, then our solution is the same as his and, generically, it is unique.¹ On the other hand, when Ross's conditions are not met, then our model can be solved while Ross's cannot. Further, we illustrate that our solution is far simpler and allows a closed-form solution that is accurate when the discount rate is close to 1.

To understand the economics of the condition $S \leq T$, consider what happens if the economy evolves in a standard multinomial tree with no upper or lower bound on the state space. For each extra time period, we get at least two new states since we can go up from highest state and down from the lowest state. Therefore, in this case $S > T$, so we see that recovery is impossible because the number of states is higher than the number of time periods. Hence, achieving recovery without further assumptions is typically impossible in most standard models of finance where the state space grows in this way. In other words, our model provides a fundamentally different way—via our simple counting argument—to understand the critique of Borovicka et al. (2016) that recovery is impossible in standard models.

Nevertheless, we show that recovery is possible even when $S > T$ under certain conditions. While maintaining that probabilities can be fully general (and, indeed, allow growth), we assume that the utility function is given via a limited number of parameters. Again, we simply need to make our counting argument work. To do this, we show that if the marginal utilities can be written as functions of N parameters, then recovery is possible as long as $N + 1 < T$. This large state-space framework is what we use empirically as discussed further below.

We illustrate how our method works in the context of three specific models, namely Mehra and Prescott (1985), Cox et al. (1979), and a simple non-Markovian economy. For each economy, we generate model-implied prices and seek to recover natural probabilities and preferences using our method. This provides an illustration of how our method works, its robustness, and its shortcomings. For Mehra and Prescott (1985), we show that $S > T$ so general recovery is impossible, but when we restrict the class of utility functions, then we achieve recovery. For the binomial Cox–Ross–Rubinstein model (the discrete-time version of Black and Scholes 1973), we show that recovery is impossible even under restrictive utility specifications because consumption growth is iid, which leads to a flat term structure, a pricing matrix of a lower rank, and a continuum of solutions for probabilities and preferences. While the former two models fall in the setting of Borovicka et al. (2016) (with a non-zero martingale

component), we also show how recovery is possible in the non-Markovian setting, which falls outside the framework of Borovicka et al. (2016) and Ross (2015), illustrating the generality of our framework in terms of the allowed probabilities.

Finally, we implement our methodology empirically using a large data set of call and put options written on the S&P 500 stock market index over the time period 1996–2015. We estimate state price densities for multiple future horizons and recover probabilities and preferences each month. Based on the recovered probabilities, we derive the risk and expected return over the future month from the physical distribution of returns using four different methods. The recovered expected returns vary substantially across specifications, challenging the empirical robustness of the results. The recovered expected returns have weak predictive power for the future realized returns, even when we exclude the global financial crisis. We can also recover ex ante volatilities, which have much stronger predictive power for future realized volatility. We note that a rejection of the recovered distribution is a rejection of the joint hypothesis of the general recovery methodology and the specific empirical choices, including the state space and the available options.

The literature on recovery theorems is quickly expanding.² Bakshi et al. (2015) and Audrino et al. (2014) empirically test the restrictions of Ross's recovery theorem. Martin and Ross (2018) apply the recovery theorem in a term structure model in which the driving state variable is a stationary Markov chain, illustrating the role played by the (infinitely) long end of the yield curve, a role already recognized in Kazemi (1992). Several papers focus on generalizing the underlying Markov process to a continuous-time process with a continuum of values and an infinite horizon (Carr and Yu 2012, Linetsky and Qin 2016) and Walden (2017) in particular derive intuitive results on the importance of recurrence. All these papers impose time-homogeneity of the underlying Markov process.³ Qin and Linetsky (2017) go beyond the Markov assumption, discussing factorization of stochastic discount factors and recovery in a general semimartingale setting.

These approaches require an infinite time horizon, while our approach only requires the observed finite-maturity data. Indeed, the martingale decomposition used by Borovicka et al. (2016) is only defined over an infinite horizon, as is the recurrence condition used by Walden (2017) and the factorization of Qin and Linetsky (2017).⁴

² Prior to Ross (2015), the dynamics of the risk-neutral density and the physical density along with the pricing kernel has been extensively researched using historical option or equity market data (e.g., Jackwerth, 2000; Jackwerth and Rubinstein, 1996; Bollerslev and Todorov, 2011; Ait-Sahalia and Lo, 2000; Rosenberg and Engle, 2002; Bliss and Panigirtzoglou, 2004 and Christoffersen et al., 2013).

³ See also Schneider and Trojani (2018) who focus on recovering moments of the physical distribution and Malamud (2016) who shows that knowledge of investor preferences is not necessarily enough to recover physical probabilities when option supply is noisy, but shows how recovery can be feasible when the volatility of option supply shocks is also known.

⁴ Said differently, if we observe data from a finite number of time periods from an economy satisfying the conditions on Borovicka et al. (2016),

¹ Generically means that the result holds for all parameters except on a "small" set of parameters of zero measure. For the measure-zero set of parameters where a certain matrix of prices has less than full rank such that there is a continuum of solutions to our generalized recovery problem, we show that the multi-period version of Ross's problem also has a continuum of solutions.

Our paper contributes to the literature by characterizing recovery of any probability distributions observed over a finite number of periods, by proving a simple solution and its closed-form approximation, and by providing natural empirical tests of our generalized method. Rather than relying on specific probabilistic assumptions (Markov processes and ergodicity) as in Ross (2015) and Borovicka et al. (2016), we follow the tradition of general equilibrium (GE) theory, where Debreu (1970) pioneered the use of Sard's theorem and differential topology. Bringing Sard's theorem into the recovery debate provides new economic insight on when recovery is possible.⁵ Indeed, the martingale decomposition applied by Borovicka et al. (2016) relies on knowing the infinite-time distribution of Markov processes, which imposes much more structure than needed and removes the focus from the essence of the recovery problem, namely the number of economic variables versus economic restrictions.

2. Ross's recovery theorem

This section briefly describes the mechanics of the recovery theorem of Ross (2015) as a background for understanding our generalized result in which we relax the assumption that transition probabilities are time-homogeneous.

The idea of the recovery theorem is most easily understood in a one-period setting. In each time period 0 and 1, the economy can be in a finite number of states which we label $1, \dots, S$. Starting in any state i , there exists a full set of Arrow–Debreu securities, each of which pays 1 if the economy is in state j at date 1. The price of these securities is given by $\pi^{i,j}$.

The objective of the recovery theorem is to use information about these observed state prices to infer physical probabilities $p^{i,j}$ of transitioning from state i to j . We can express the connection between Arrow–Debreu prices and physical probabilities by introducing a pricing kernel m such that for any $i, j = 1, \dots, S$,

$$\pi^{i,j} = p^{i,j} m^{i,j}. \tag{1}$$

It takes no more than a simple one-period binomial model to convince oneself that if we know the Arrow–Debreu prices in one and only one state at date 0, then there is in general no hope of recovering physical probabilities. In short, we cannot separate the contribution to the observed Arrow–Debreu prices from the physical probabilities and the pricing kernel.

The key insight of the recovery theorem is that by assuming that we know the Arrow–Debreu prices for all the possible starting states, then with additional structure on the pricing kernel, we can recover physical probabilities. We note that knowing the prices in states we are not currently in (“parallel universes”) is a strong assumption.

then there is no unique Markov decomposition. Recurrence means that each state is being visited infinitely often, so it can only be defined over an infinite horizon. The factorization of Qin and Linetsky (2017) relies on limits of T -forward measures, as T goes to infinity.

⁵ We thank Steve Ross for pointing out the historical role of Sard's theorem in general equilibrium theory.

In any event, under this assumption, Ross's result is that there exists a unique set of physical probabilities $p^{i,j}$ for all $i, j = 1, \dots, S$ such that Eq. (1) holds if the matrix of Arrow–Debreu prices is irreducible and if the pricing kernel m has the form known from the standard representative agent models:

$$m^{i,j} = \delta \frac{u^j}{u^i}, \tag{2}$$

where $\delta > 0$ is the discount rate, and $u = (u^1, \dots, u^S)$ is a vector with strictly positive elements representing marginal utilities.

The proof can be found in Ross (2015), but here we note that counting equations and unknowns certainly makes it plausible that the theorem is true: there are S^2 observed Arrow–Debreu prices and hence S^2 equations. Because probabilities from a fixed starting state sum to one, there are $S(S - 1)$ physical probabilities. It is clear that scaling the vector u by a constant does not change the equations, and thus we can assume that $u^1 = 1$ so that u contributes with an additional $S - 1$ unknowns. Adding to this the unknown δ leaves us exactly with a total of S^2 unknowns. The fact that there is a unique strictly positive solution hinges on the Frobenius theorem for positive matrices.

It is important in Ross's setting, as it will be in ours that a state corresponds to a particular level of the marginal utility of consumption. This level does not depend on calendar time. In our empirical implementation, a state will correspond to a particular level of the S&P 500 index.

The most troubling assumption, however, in the theorem above is that we must also know state prices from starting states that we are currently not in. It is hard to imagine data that would allow us to know these in practice. Ross's way around this assumption is to leave the one-period setting and assume that we have information about Arrow–Debreu prices from several future periods, and then use a time-homogeneity assumption to recover the same information that we would be able to obtain from the equations above.

We therefore consider a discrete-time economy with time indexed by t , states indexed by $s = 1, \dots, S$, and $\pi_{t,t+\tau}^{i,j}$ denoting the time- t price in state i of an Arrow–Debreu security that pays 1 in state j at date $t + \tau$. The multi-period analogue of Eq. (1) becomes

$$\pi_{t,t+\tau}^{i,j} = p_{t,t+\tau}^{i,j} m_{t,t+\tau}^{i,j}. \tag{3}$$

Similarly, the multi-period analogue to Eq. (2) is the following assumption, which again follows from the existence of a representative agent with time-separable utility:

Assumption 1 (Time-separable utility). There exists a $\delta \in (0, 1]$ and marginal utilities $u^j > 0$ for each state j such that for all times τ , the pricing kernel can be written as

$$m_{t,t+\tau}^{i,j} = \delta^\tau \frac{u^j}{u^i}. \tag{4}$$

Critically, to move to a multi-period setting, Ross makes the following additional assumption of time-homogeneity to implement his approach empirically:

Assumption 2 (Time-homogeneous probabilities). For all states i, j , and time horizons $\tau > 0$, $p_{t,t+\tau}^{i,j}$ does not depend on t .

This assumption is strong and not likely to be satisfied empirically. We note that Assumptions 1 and 2 together imply that risk neutral probabilities are also time-homogeneous, a prediction that can also be rejected in the data.

In this paper, we dispense with the time-homogeneity Assumption 2. We start by maintaining Assumption 1 but later consider a broader assumption that can be used in a large state space.

3. A generalized recovery theorem

The assumption of time-separable utility is consistent with many standard models of asset pricing, but the assumption of time-homogeneity is much more troubling. It restricts us from working with a growing state space (as in standard binomial models), and it makes numerical implementation extremely hard and non-robust, because trying to fit observed state prices to a time-homogeneous model is extremely difficult. Furthermore, the main goal of the recovery exercise is to recover physical transition probabilities from the current states to all future states over different time horizons. Insisting that these transition probabilities arise from constant one-period transition probabilities is a strong restriction. We show in this section that by relaxing the assumption of time-homogeneity of physical transition probabilities, we can obtain a problem that is easier to solve numerically and that allows for a much richer modeling structure. We show that our extension contains the time-homogeneous case as a special case and therefore ultimately should allow us to test whether the time-homogeneity assumption can be defended empirically.

3.1. A Noah’s ark example: two states and two dates

To get the intuition of our approach, we start by considering the simplest possible case with two states and two time periods. Consider the simple case in which the economy has two possible states (1, 2) and two time periods starting at time t and ending on dates $t + 1$ and $t + 2$. If the current state of the world is state 1, then Eq. (3) consists of four equations:

$$\begin{aligned} \pi_{t,t+1}^{1,1} &= p_{t,t+1}^{1,1} & m_{t,t+1}^{1,1} \\ \pi_{t,t+1}^{1,2} &= (1 - p_{t,t+1}^{1,1}) & m_{t,t+1}^{1,2} \\ \pi_{t,t+2}^{1,1} &= p_{t,t+2}^{1,1} & m_{t,t+2}^{1,1} \\ \pi_{t,t+2}^{1,2} &= (1 - p_{t,t+2}^{1,1}) & m_{t,t+2}^{1,2} \end{aligned} \tag{5}$$

2 unknowns
4 unknowns

We see that we have 4 equations with 6 unknowns, so this system cannot be solved in full generality. However, the number of unknowns is reduced under the assumption of time-separable utility (Assumption 1). To see that

most simply, we introduce the notation h for the normalized vector of marginal utilities:

$$h = \left(1, \frac{u^2}{u^1}, \dots, \frac{u^S}{u^1} \right)' \equiv (1, h_2, \dots, h_S)', \tag{6}$$

where we normalize by u^1 . With this notation and the assumption of time-separable utility, we can rewrite the system (5) as follows:

$$\begin{aligned} \pi_{t,t+1}^{1,1} &= p_{t,t+1}^{1,1} \delta \\ \pi_{t,t+1}^{1,2} &= (1 - p_{t,t+1}^{1,1}) \delta h_2 \\ \pi_{t,t+2}^{1,1} &= p_{t,t+2}^{1,1} \delta^2 \\ \pi_{t,t+2}^{1,2} &= (1 - p_{t,t+2}^{1,1}) \delta^2 h_2. \end{aligned} \tag{7}$$

This system now has 4 equations with 4 unknowns, so there is hope to recover the physical probabilities p , the discount rate δ , and the ratio of marginal utilities h .

Before we proceed to the general case, it is useful to see how the problem is solved in this case. Moving h_2 to the left side and adding the first two and the last two equations gives us two new equations

$$\begin{aligned} \pi_{t,t+1}^{1,1} + \pi_{t,t+1}^{1,2} \frac{1}{h_2} - \delta &= 0 \\ \pi_{t,t+2}^{1,1} + \pi_{t,t+2}^{1,2} \frac{1}{h_2} - \delta^2 &= 0. \end{aligned} \tag{8}$$

Solving Eq. (8) for h_2 yields $\frac{1}{h_2} = (\delta - \pi_{t,t+1}^{1,1}) / \pi_{t,t+1}^{1,2}$, and we can further arrive at

$$\pi_{t,t+2}^{1,1} - \frac{\pi_{t,t+2}^{1,2} \pi_{t,t+1}^{1,1}}{\pi_{t,t+1}^{1,2}} + \frac{\pi_{t,t+2}^{1,2}}{\pi_{t,t+1}^{1,2}} \delta - \delta^2 = 0. \tag{9}$$

Hence, we can solve the two-state model by (i) finding δ as a root of the second degree polynomial (9), (ii) computing the marginal utility ratio h_2 from Eq. (8); and (iii) computing the physical probabilities by rearranging Eq. (7).

3.2. General case: notation

Turning to the general case, recall that there are S states and T time periods. Without loss of generality, we assume that the economy starts at date 0 in state 1. This allows us to introduce some simplifying notation since we do not need to keep track of the starting time or the starting state – we only need to indicate the final state and time horizon over which we are considering a specific transition.

Accordingly, let $\pi_{\tau s}$ denote the price of receiving 1 at date τ if the realized state is s and collect the set of observed state prices in a $T \times S$ matrix Π defined as

$$\Pi = \begin{bmatrix} \pi_{11} & \dots & \pi_{1S} \\ \vdots & & \vdots \\ \pi_{T1} & \dots & \pi_{TS} \end{bmatrix}. \tag{10}$$

Similarly, letting $p_{\tau s}$ denote the physical transition probabilities of going from the current state 1 to state s in τ periods, we define a $T \times S$ matrix P of physical probabilities. Note that $p_{\tau s}$ is *not* the probability of going from state τ to s (as in the setting of Ross 2015), but rather the first index denotes time for the purpose of the derivation of our theorem.

From the vector of normalized marginal utilities h defined as in Eq. (6), we define the S -dimensional diagonal matrix $H = \text{diag}(h)$. Further, we construct a T -dimensional diagonal matrix of discount factors as $D = \text{diag}(\delta, \delta^2, \dots, \delta^T)$.

3.3. Generalized recovery

With this notation in place, the fundamental TS equations linking state prices and physical probabilities, assuming utilities depend on current state only, can be expressed in matrix form as

$$\Pi = DPH. \tag{11}$$

Note that the (invertible) diagonal matrices H and D depend only on the vector h and the constant δ , so if we can determine these, we can find the matrix of physical transition probabilities from the observed state prices in Π :

$$P = D^{-1}\Pi H^{-1}. \tag{12}$$

Since probabilities add up to 1, we can write $Pe = e$, where $e = (1, \dots, 1)'$ is a vector of ones. Using this identity, we can simplify Eq. (12) such that it only depends on δ and h :

$$\Pi H^{-1}e = DPe = De = (\delta, \delta^2, \dots, \delta^T)'. \tag{13}$$

To further manipulate this equation it will be convenient to work with a division of Π into block matrices:

$$\Pi = \begin{bmatrix} \Pi_1 & \Pi_2 \end{bmatrix} = \begin{bmatrix} \Pi_{11} & \Pi_{12} \\ \Pi_{21} & \Pi_{22} \end{bmatrix}. \tag{14}$$

Here, Π_1 is a column vector of dimension T , where the first $S - 1$ elements are denoted by Π_{11} , and the rest of the vector is denoted Π_{21} . Similarly, Π_2 is a $T \times (S - 1)$ matrix, where the first $S - 1$ rows are called Π_{12} , and the last rows are called Π_{22} . With this notation and the fact that $H(1, 1) = h(1) = 1$, we can write Eq. (13) as

$$\Pi_1 + \Pi_2 \begin{bmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{bmatrix} = \begin{bmatrix} \delta \\ \vdots \\ \delta^T \end{bmatrix}, \tag{15}$$

where of course $h_s^{-1} = \frac{1}{h_s}$. Given that these equations are linear in the inverse marginal utilities h_s^{-1} , it is tempting to solve for these. To solve for these $S - 1$ marginal utilities, we consider the first $S - 1$ equations

$$\Pi_{11} + \Pi_{12} \begin{bmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{bmatrix} = \begin{bmatrix} \delta \\ \vdots \\ \delta^{S-1} \end{bmatrix}, \tag{16}$$

with solution⁶

$$\begin{bmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{bmatrix} = \Pi_{12}^{-1} \left(\begin{bmatrix} \delta \\ \vdots \\ \delta^{S-1} \end{bmatrix} - \begin{bmatrix} \pi_{11} \\ \vdots \\ \pi_{S-1,1} \end{bmatrix} \right). \tag{17}$$

Hence, if δ were known, we would be done. Since δ is a discount rate, it is reasonable to assume that it is close to

⁶ Of course, to invert Π_{12} it must have full rank. As long as Π_2 has full rank, we can reorder the rows to ensure that Π_{12} also has full rank.

one over short time periods. We later use this insight to derive a closed-form approximation that is accurate as long as we have a reasonable sense of the size of δ . For now, we proceed for general unknown δ .

We thus have the utility ratios given as a linear function of powers of δ . The remaining $T - S + 1$ equations give us

$$\Pi_{21} + \Pi_{22} \begin{bmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{bmatrix} = \begin{bmatrix} \delta^S \\ \vdots \\ \delta^T \end{bmatrix}, \tag{18}$$

and from this we see that if we plug in the expression for the utility ratios found above, we end up with $T - S + 1$ equations, each of which involves a polynomial in δ of degree at most T . If $T = S$, then δ is a root to a single polynomial, so at most a finite number of solutions exist. If $T > S$, then generically no solution exists for general Arrow–Debreu prices Π since δ must simultaneously solve several polynomial equations (where “generically” means almost surely as defined just below Proposition 1). However, if the prices are generated by the model, then a solution exists and it will almost surely be unique. To be precise, we say that Π has been “generated by the model” if there exist δ , P , and H such that Π can be found from the right-hand side of Eq. (11). The following theorem formalizes these insights (using Sard’s theorem):

Proposition 1 (Generalized recovery). Consider an economy satisfying Assumption 1 with Arrow–Debreu prices for each of the T time periods and S states. The recovery problem has

1. a continuum of solutions if $S > T$;
2. at most S solutions if the submatrix Π_2 has full rank and $S = T$;
3. no solution generically in terms of an arbitrary positive matrix Π and $S < T$;
4. a unique solution generically if Π has been generated by the model and $S < T$.

The proof of this and all following propositions are in the appendix. The proposition states our results using the notion “generically,” which means that they fail to hold at most for a set of measure zero. Said differently, if someone picks parameters “at random,” then our results hold almost surely.⁷

Further, since Sard’s theorem is not a standard tool in asset pricing theory, some words here on the basic intuition behind our use of the theorem are in order. To get started, consider a linear function $f(x) = Ax$ from R^m to R^n given by the $n \times m$ matrix A . We know that if $n = m$ and A has full rank, then the image of A is all of R^n , i.e., every point of R^n is being “hit” by A . If, however, $n > m$, then the image of A is a linear subspace of R^n , which is vanishingly small (has Lebesgue measure 0 in R^n). By Sard’s theorem, we can extend this result to a nonlinear smooth function f and still conclude that when $n > m$, the image of f is vanishingly small. Said differently, there exists no solution x to $f(x) = y$ generically (i.e., if you pick a random y , then

⁷ We note that the fact that our results hold only generically is not a consequence of our solution method—indeed, there exist counterexamples for special sets of parameters as discussed in our examples.

almost surely no solution exists).⁸ While this result is relatively simple, it has wide-ranging implications and, indeed, allows us to prove Proposition 1.

4. Generalized recovery versus other forms of recovery

Proposition 1 provides a simple way to understand when recovery is possible, namely, essentially when the number of time periods T is at least as large as the number of states S . We now show how our method relates to Ross’s method and other recovery results.

4.1. Generalized recovery in a Ross economy

We first show that our method generalizes Ross’s recovery method in the sense that if we are in a Ross economy, then any solution to Ross’s problem has a corresponding solution to our problem.

It is important to be clear about the terminology here. In Ross’s recovery problem, physical transition probabilities are specified in terms of a one-period transition probability matrix \bar{P} that includes transition probabilities from states that we are currently not in (“parallel universes”). Our problem focuses on recovering the matrix P of multi-period transition probabilities as seen from the state we are in at time 0, which we take to be state 1. We say that P is generated from \bar{P} if the k th row of P is equal to the first row of \bar{P}^k . The same terminology can be applied to state prices, of course.⁹

Proposition 2 (Generalized recovery works in a Ross economy). If observed prices Π over $S = T$ time periods are generated by a Ross economy (i.e., an irreducible matrix $\bar{\Pi}$ of one-period state prices and probabilities P satisfying Assumptions 1 and 2), then

1. The matrix P generated from \bar{P} is a solution to our generalized recovery problem.
2. P is a unique solution to our generalized recovery problem generically in the space of Ross price matrices $\bar{\Pi}$.
3. If Π_{12} has full rank, then Ross’s parallel universe prices $\bar{\Pi}$ can be derived uniquely from multi-period prices Π observed from the current state. Otherwise, there can exist a continuum of Ross prices $\bar{\Pi}$ consistent with the observed prices. The rank condition is satisfied generically in the space of Ross price matrices.

Part 1 of the proposition confirms that any solution to Ross’s recovery problem corresponds to a solution to our generalized problem. Part 2 of the proposition considers the deeper question of uniqueness. Ross establishes

⁸ On a more technical note, Sard’s theorem in fact states that if M is the set of critical points of f (i.e., the set of points for which the Jacobian matrix of f has rank strictly smaller than n), then $f(M)$ has Lebesgue measure zero in R^n . When $n > m$, all points are critical points, and therefore in this case $f(M)$ is the same as the image of f , which is what we need for our proof.

⁹ The notion of generating P from \bar{P} is based on the fact that in a Ross economy, the matrix of probabilities of going from state i to state j in k time periods is given by \bar{P}^k . Likewise, the k -period state prices are given by $\bar{\Pi}^k$.

a unique solution while our generalized recovery solution in our earlier Proposition 1 only narrows the solution set down to at most $S = T$ solutions. Interestingly, Proposition 2 shows that our method too yields a unique solution when prices come from a Ross economy, generically. Thus, in this sense, nothing is lost by using generalized recovery even when we are in a Ross economy.

One way to understand this result is to note that Ross’s problem comes down to solving a characteristic polynomial, and similarly, our generalized recovery problem can be solved via the polynomial given by Eq. (18). Even though these polynomials come from different sets of equations, it turns out that they have the same roots when Ross’s assumptions are satisfied.

Finally, part 3 of the proposition deals with the issue that some of our results only hold “generically,” that is, for almost all parameters. One might ask whether Ross also has a similar problem for the (small set of) remaining parameters. The answer turns out to be yes and for a reason that has not yet been discussed in the context of Ross’s method. The issue is that Ross finds a unique solution given his parallel universe price matrix $\bar{\Pi}$, but where does this matrix come from? In any real-world application, we start with observed prices Π over time as in our generalized recovery setting. When Ross implements his model empirically, he must first find his $\bar{\Pi}$ from the observed Π , and then use his recovery method (but he does not consider the mathematics of the first step, getting $\bar{\Pi}$ from Π). Part 3 of the proposition shows that Ross has the same problem as we do for the small set of parameters, where Π_{12} has less than full rank. In other words, his lack of uniqueness arises from the difficulty in finding the price matrix $\bar{\Pi}$. Interestingly, this may have been unnoticed since Ross takes $\bar{\Pi}$ as given in his theoretical analysis (and shows that his recovery is unique for each $\bar{\Pi}$).

This last point is most clearly seen through an example. Consider two different one-period transition probability matrices that are both irreducible:

$$\bar{P} = \begin{pmatrix} 1 & 1 & 1 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix} \text{ and } \bar{P}' = \begin{pmatrix} 1 & 1 & 1 \\ \frac{1}{3} & \frac{2}{3} & 0 \\ \frac{1}{3} & 0 & \frac{2}{3} \end{pmatrix}.$$

If we assume that the current state is state 1, then since all powers of the matrices \bar{P} and \bar{P}' have the same first row, namely $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, it follows that the matrices P and P' (i.e., the physical transition probabilities as seen from state 1) generated by \bar{P} and \bar{P}' become the same matrix

$$P = P' = \begin{pmatrix} 1 & 1 & 1 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

For given discount factors D and marginal utilities H , $\Pi = DP^H$ and $\Pi' = DP'^H$ are then the same, and hence observing the 3×3 matrix of state prices Π would not allow us to distinguish between the physical transition matrices \bar{P} and \bar{P}' . The problem is not mitigated by observing

more periods. It is simply impossible in a world where we cannot observe parallel universe prices to distinguish between the two irreducible matrices. In our approach, we do not seek to recover the one-period transition probabilities. Rather, we recover the matrix P , and our ability to do so depends on the rank of a submatrix the Π matrix. For example, if we let $\delta = 0.98$, and let $h_1 = 1, h_2 = 0.9, h_3 = 0.8$, then the submatrix of state prices Π_{12} has rank 1, and this means that we would not have unique recovery either.

4.2. Ross recovery in our generalized economy

We now establish that our formulation is *strictly* more general by showing that for many typical price matrices (e.g., those observed in the data), no solution exists for Ross's recovery problem, even though a solution exists for the generalized recovery problem.

Proposition 3 (Generalized recovery is more general). With $S = T$, there exists set of parameters with positive Lebesgue measure for the generalized recovery problem where no solution exists for Ross's recovery problem. With $S < T$, generically among price matrices for the generalized recovery problem, there exists no solution to Ross's recovery problem.

This proposition shows that generalized recovery can be useful because it can match a broader class of market prices in addition to the basic advantage that it starts with the observed multi-period prices (rather than parallel universe prices).

4.3. Recovery in infinite horizon

In addition to generalizing Ross's method, our result also provides a simple and intuitive way of understanding why, for example, growth may present a challenge for recovery (cf. the critique of Borovicka et al. 2016, that recovery is infeasible in standard models). Indeed, we provide a simple counting argument: Suppose that the economy has growth such that for each extra time period, the economy can increase from the previously highest state and go down from the previously lowest state. Then we get two new states for each new time period, which implies that $S > T$ such that recovery is impossible. Nevertheless, we can still achieve recovery in such a large state space if we consider a class of pricing kernels that is sufficiently low-dimensional, as we discuss below in Section 6.

Our argument is very different from that of Borovicka et al. (2016) who rely on a martingale decomposition, which requires an infinite time horizon. Our counting argument is simple and is based on a finite horizon, consistent with the data observed in practice.

Our finite-horizon recovery theorem is therefore also markedly distinct from the existing approaches that exist in continuous-time models in that we make no reference to, and have no need for, recurrence or stationarity conditions. In a diffusion setting, Walden (2017) shows the fundamental role of recurrence as a necessary condition for recovery in these models. Recurrence essentially means that each state is being visited infinitely often, so it can only be defined over an infinite horizon. Recurrence

bears some resemblance to Ross's condition of irreducibility in that an infinite time extension of an irreducible chain would be recurrent. The result of Walden (2017) is intuitive since, when states are visited infinitely often, we have a chance to recover probabilities.

Our approach can naturally be used to consider whether recovery is possible in a finite-time version of an infinite-horizon process (i.e., even if a process is defined over an infinite horizon, we can ask what happens if we only see it over a couple of years). Further, we can show via some examples that recovery may even be possible for nonrecurrent processes or processes with growth.

To give a simple example of this, consider a two-period nonhomogeneous Markov process with two states defined from the probability transition matrices for each time

$$\bar{P}(0, 1) = \begin{pmatrix} 0.4 & 0.6 \\ 0.5 & 0.5 \end{pmatrix} \text{ and, for } t \geq 1, \\ \bar{P}(t, t+1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

In the first period, the process either stays in its current state or jumps to the other state, but after that, the process is absorbed in its current state. If we only observe prices for two time periods, then this is clearly the restriction of a nonrecurrent process. Given that $S = T = 2$, our counting argument shows that generalized recovery is feasible.

We could also imagine a process with growth, starting in the "lowest state" 1 and evolving according to a transition matrix specified as an upward drifting process. To give a simple illustration, imagine Assumptions 1 and 2 hold and that the one-period transition matrix of physical probabilities across five states is given as

$$\bar{P} = \begin{pmatrix} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.1 & 0.5 & 0.4 & 0 & 0 \\ 0 & 0.1 & 0.5 & 0.4 & 0 \\ 0 & 0 & 0.1 & 0.5 & 0.4 \\ 0 & 0 & 0 & 0.5 & 0.5 \end{pmatrix}.$$

If we observe prices over five time periods, then our counting argument is satisfied $S = T = 5$, and we see that it is not growth per se that makes recovery impossible—it is the expanding state space necessary to accommodate models with growth that may cause problems.

In summary, our results complement those in the literature in two ways. First, generalized recovery may work when other methods do not and vice versa. Second, generalized recovery provides an economic intuition in finite economies, while other methods do so in infinite-horizon economies.

4.4. Flat term structure and risk neutrality

We finally note that the very special case of an observed flat term structure of interest rates has some special properties. In particular, with a flat term structure there exists a solution to the problem in which the representative agent is risk neutral, echoing an analogous result by Ross.

To see this result, note that the price of a zero-coupon bond with maturity τ is equal to the sum of the τ th row of

Π , which we write as $(\Pi e)_\tau$. Having a flat term structure means that the yield on the zero-coupon bonds does not depend on maturity, i.e., that there exists a constant r such that

$$\frac{1}{(1+r)^\tau} = (\Pi e)_\tau. \tag{19}$$

Let the $T \times S$ matrix Q contain the risk-neutral transition probabilities seen from the starting state, i.e., the k th row of Q gives us the risk-neutral probabilities of ending in the different states at date k .

Proposition 4 (Flat term structure). *Suppose that the term structure of interest rates is flat, i.e., there exists $r > 0$ such that $\frac{1}{(1+r)^\tau} = (\Pi e)_\tau$ for all $\tau = 1, \dots, T$. Then the recovery problem is solved with equal physical and risk-neutral probabilities, $P = Q$. This means that either the representative agent is risk neutral or the recovery problem has multiple solutions.*

We note that this result should be interpreted with caution. The knife-edge (i.e., measure zero) case of a flat term structure may well be generated by the knife-edge case of a price matrix Π with low rank, which implies that a continuum of solutions may exist, and the representative agent may well be risk averse (as one would expect). Intuitively, a flat term structure may be generated by a Π with so much symmetry that it has a low rank.

5. Closed-form recovery

The recovery problem is almost linear, except for the powers of the discount rate δ that enter into the problem as a polynomial. In practical implementations over the time horizons where options are liquid, a linear approximation provides an accurate approximation given that δ is known to be close to one at an annual horizon.

The linear approximation is straightforward. To linearize the discounting of δ^τ around a point δ_0 (say, $\delta_0 = 0.97$), we write $\delta^\tau \approx a_\tau + b_\tau \delta$ for known constants a_τ and b_τ . Based on the Taylor expansion $\delta^\tau \approx \delta_0^\tau + \tau \delta_0^{\tau-1} (\delta - \delta_0)$, we have $a_\tau = -(\tau - 1)\delta_0^\tau$ and $b_\tau = \tau \delta_0^{\tau-1}$. As seen in Fig. 2, the approximation is accurate for $\delta \in [0.94, 1]$ for time horizons less than two years.

With the linearization of the polynomials in δ , the equations for the recovery problem (13) become the following:


$$\begin{pmatrix} \pi_{11} \\ \vdots \\ \pi_{T1} \end{pmatrix} + \begin{pmatrix} \pi_{12} & \dots & \pi_{1S} \\ \vdots & & \vdots \\ \pi_{T2} & \dots & \pi_{TS} \end{pmatrix} \begin{pmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{pmatrix} = \begin{pmatrix} a_1 + b_1 \delta \\ \vdots \\ a_T + b_T \delta \end{pmatrix}. \tag{20}$$

which we can rewrite as a system of T equations in S unknowns as


$$\begin{pmatrix} -b_1 & \pi_{12} & \dots & \pi_{1S} \\ \vdots & \vdots & & \vdots \\ -b_T & \pi_{T2} & \dots & \pi_{TS} \end{pmatrix} \begin{pmatrix} \delta \\ h_2^{-1} \\ \vdots \\ h_S^{-1} \end{pmatrix} = \begin{pmatrix} a_1 - \pi_{11} \\ \vdots \\ a_T - \pi_{T1} \end{pmatrix}. \tag{21}$$

Rewriting this equation in matrix form as

$$Bh_\delta = a - \pi_1 \tag{22}$$



Panel A: $t = 2$ years



Panel B: $t = 0.5$ years

Fig. 2. Closed-form solution: approximation error. The figure shows that the generalized recovery problem is very close to being linear. We show that the only nonlinearity comes from the discount rate δ due to the powers of time, δ^t . However, the function $\delta \rightarrow \delta^t$ is very close to being linear for the relevant range of annual discount rates, say $\delta \in [0.94, 1]$, and the relevant time periods that we study. Panel A plots the discount function and the linear approximation around $\delta_0 = 0.97$ given a horizon of $t = 2$ years. Panel B plots the same for a horizon of a half year.

we immediately see the closed-form solution

$$h_\delta = \begin{cases} B^{-1}(a - \pi_1) & \text{for } S = T \\ (B'B)^{-1}B'(a - \pi_1) & \text{for } S < T. \end{cases} \tag{23}$$

We see that when $S = T$, we simply need to solve S linear equations with S unknowns. When $S < T$, we could simply just consider S equations and ignore the remaining $T - S$ equations.

More broadly, if $S < T$ and we start with prices Π that are not exactly generated by the model (e.g., because of noise in the data), then Eq. (23) provides the values of δ and the vector h that best approximate a solution in the sense of least squares.

The following theorem shows that the closed-form solution is accurate as long as the value of δ_0 is close to the true discount rate:

Proposition 5 (Closed-form solution). *If prices are generated by the model and B has full rank $S \leq T$, then the closed-form solution (23) approximates the true solution in the following sense: the distance between the true solution $(\bar{\delta}, \bar{h}, \bar{P})$ and the approximate solution (δ, h, P) approaches 0 faster than $(\delta_0 - \bar{\delta})$ as δ_0 approaches $\bar{\delta}$.*

6. Recovery in a large state space

A challenge in implementing the Ross recovery theorem is that it does not allow for an expanding set of states as we know it, for example, from binomial models and multinomial models of option pricing. Simply stated, the expanding state space in a binomial model adds more unknowns for each time period than equations even under the assumption of utility functions that depend on the current state only. We next show how we handle an expanding state space in our model.

We have in mind a case in which the number of states S is larger than the number of time periods T . In a standard binomial model, for example, with two time periods we need five states corresponding to the different values that the stock can take over its path. The key to solving this problem is to reduce the dimensionality of the utility ratios captured in the vector h . To do that, we replace Assumption 1 with the following assumption that the pricing kernels belong to a parametric family with limited dimensionality.

*Assumption 1** (General utility with N parameters). The pricing kernel at time τ in state s (given the initial state 1 at time 0) can be written as

$$m_{0,\tau}^{1,s} = \delta^\tau h_s(\theta), \tag{24}$$

where $\delta \in (0, 1]$ and $h(\cdot) > 0$ is a one-to-one C^∞ smooth function of the parameter $\theta \in \Theta$, an embedding from $\Theta \subset \mathbb{R}^N$ to \mathbb{R}^S , and Θ has a non-empty interior.

With a large number of unknowns compared to the number of equations, we need to restrict the set of unknowns, and this is done by assuming that the utilities are parameterized by a lower-dimensional set Θ .

6.1. A large discrete state space

Let us first consider two simple examples of how we can parameterize marginal utilities with a low-dimensional set of parameters. First, we consider a simple linear expression for the marginal utilities, and then we discuss the case of constant relative risk aversion (a nonlinear mapping from risk aversion parameters Θ to marginal utilities).

We start with a simple linear example of how the parameterization works. We consider a matrix B of full rank and dimension $(S - 1) \times N$ such that

$$\begin{bmatrix} h_2^{-1} \\ \vdots \\ h_S^{-1} \end{bmatrix} = \begin{pmatrix} a_1 \\ \vdots \\ a_{S-1} \end{pmatrix} + \begin{pmatrix} b_{11} & \dots & b_{1N} \\ \vdots & & \vdots \\ b_{S-1,1} & \dots & b_{S-1,N} \end{pmatrix} \begin{bmatrix} \theta_1 \\ \vdots \\ \theta_N \end{bmatrix} = A + B\theta. \tag{25}$$

Combining this equation with the recovery problem (15) gives

$$(\Pi_1 + \Pi_2 A) + \Pi_2 B \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_N \end{pmatrix} = \begin{pmatrix} \delta \\ \vdots \\ \delta^T \end{pmatrix}. \tag{26}$$

This equation has *exactly* the same form as our original recovery problem (15), but now $\Pi_1 + \Pi_2 A$ plays the role

of Π_1 , similarly $\Pi_2 B$ plays the role of Π_2 , and θ plays the role of $(h_2^{-1}, \dots, h_S^{-1})'$. The only difference is that the dimension of the unknown parameter has been reduced from $S - 1$ to N . Therefore, Proposition 1 holds as stated with S replaced by $N + 1$.

Hence, while before we could achieve recovery if $S \leq T$, now we can achieve recovery as long as $N + 1 \leq T$. In other words, recovery is possible as long as the representative agent's utility function can be specified by a number of parameters that is small relative to the number of time periods for which we have price data.

Assumption 1* also allows for the marginal utilities to be non-linear function of the risk aversion parameters θ . This generality is useful because standard utility functions can give rise to such a nonlinearity. As a simple example, consider an economy with a representative agent with CRRA preferences. In this economy, the pricing kernel in state s at time τ (given the current state 1 at time 0) is

$$m_{0,\tau}^{1,s} = \delta^\tau \left(\frac{c_s}{c_1} \right)^{-\theta}, \tag{27}$$

where c_s is the known consumption in state s of the representative agent, and θ is the unknown risk aversion parameter. Hence, Assumption 1* is clearly satisfied with $h_s^{-1}(\theta) = (c_s/c_1)^\theta$. Our generalized recovery result extends to the large state space as stated in the following proposition.

Proposition 6 (Generalized recovery in a large state space). Consider an economy satisfying Assumption 1* with Arrow–Debreu prices for each of the T time periods and S states such that $N + 1 < T$. The recovery problem has

1. no solution generically in terms of an arbitrary Π matrix of positive elements;
2. a unique solution generically if Π has been generated by the model.

As one simple application of the proposition, we can recover preferences from state prices if we know that the pricing kernel is bounded and there are sufficiently many time periods as seen in the following corollary. Said differently, using a simplified or winsorized pricing kernel (or state space) is a special case of Proposition 6.

Corollary 7 (Generalized recovery with bounded kernel). Suppose that the pricing kernel is bounded in the sense that there exist states $\bar{s} > \underline{s}$ such that $h_s = h_{\bar{s}}$ for $s > \bar{s}$ and $h_s = h_{\underline{s}}$ for $s < \underline{s}$. Then the conclusion of Proposition 6 applies, where N is the number of states from \underline{s} to \bar{s} .

6.2. Continuous state space

Finally, we note that our framework also easily extends to a continuous state space under Assumption 1* in discrete time (see Walden 2017 for the case of continuous time and continuous state space). We start with a continuous state-space density $\pi_\tau(s)$ at each time point $\tau = 1, \dots, T$ (given the current state at time 0). As before, $\pi_\tau(s)$ represents Arrow–Debreu prices or, more precisely, $\pi_\tau(s) ds$ represents the current value of receiving 1 at time τ if the state is in a small interval ds around s . Similarly, we let $p_\tau(s)$ denote the physical probability density of transitioning to s in τ periods. The fundamental recovery equations

now become

$$\pi_\tau(s) = \delta^\tau h_s(\theta) p_\tau(s). \tag{28}$$

By moving h to the left-hand side and integrating, we can eliminate the natural probabilities as before.

$$\int \pi_\tau(s) h_s^{-1}(\theta) ds = \delta^\tau. \tag{29}$$

For each time period τ , this gives an equation to help us recover the $N + 1$ unknowns, namely the discount rate δ and the parameters $\theta \in \mathbb{R}^N$. Hence, we are in the same situation as in the discrete-state model of Section 6.1, and we have recovery if there are enough time periods as stated in Proposition 6.

As before, the linear case is particularly simple. Suppose that the marginal utilities can be written as¹⁰

$$h_s^{-1}(\theta) = A(s) + B(s)\theta, \tag{30}$$

where for each s , $A(s)$ is a known scalar and $B(s)$ is a known row vector of dimension N . Using this expression, we can rewrite Eq. (29) as a simple equation of the same form as our original recovery problem (15):

$$\pi_\tau^A + \pi_\tau^B \theta = \delta^\tau, \tag{31}$$

where $\pi_\tau^A = \int \pi_\tau(s) A(s) ds$ and $\pi_\tau^B = \int \pi_\tau(s) B(s) ds$. Hence, as before, we have T equations that are linear except for the powers of the discount rate.

7. Recovery in specific models: examples

In this section we investigate recovery of specific models of interest. In a controlled environment, we show when for given state prices, our model recovers the true underlying risk-aversion parameter and time-preference parameter, along with the true multiperiod physical probabilities.

7.1. Recovery in the Mehra and Prescott (1985) model

The Mehra and Prescott (1985) model works as follows. The aggregate consumption either grows at rate $u = 1.054$ or shrinks at rate $d = 0.982$ over the next period. This consumption growth between time $t - 1$ and t is captured by a process X_t . The aggregate consumption process can be written as

$$Y_t = \prod_{s=1}^t X_s \tag{32}$$

where the initial consumption is normalized as $Y_0 = 1$.

Consumption growth X_t is a Markov process with two states, up and down. The probability of having an up state after an up state is $\phi_{uu} = Pr(X_t = u | X_{t-1} = u) = 0.43$ and, equally, the probability of staying in the down state is $\phi_{dd} = 0.43$. Hence, the probability of switching state is $\phi_{ud} = \phi_{du} = 0.57$.

The Arrow–Debreu price of receiving 1 at time t in a state $s_t = (y_t, x_t)$ is computed based on the CRRA preferences for the representative agent with risk aversion $\gamma = 4$ as

$$\pi_{0,t}^{1,s_t} = \delta^t y_t^{-\gamma} Pr(X_t = x_t, Y_t = y_t), \tag{33}$$

¹⁰ Note that $h_s^{-1}(\theta)$ denotes $\frac{1}{h_s(\theta)}$, i.e., it is *not* the inverse function of $h_s(\theta)$.

where the time-preference parameter is $\delta = 0.98$, and the physical probabilities $Pr(X_t = x_t, Y_t = y_t)$ of each state are computed based on the Markov probabilities above.¹¹

Based on this model of Mehra and Prescott (1985), we compute Arrow–Debreu prices in each state over $T = 20$ time periods and examine whether we can recover probabilities and preferences based on knowing only these prices (we have also performed the recovery for other values of T).

We first notice from Eq. (32) that consumption has growth, which immediately implies that $S > T$. This means that recovery is impossible without further assumptions. Hence, we proceed using the method concerning a large state space of Section 6. The simplest way to proceed is to assume that we know the form of the pricing kernel (33), but we do not know the risk aversion γ , the discount rate δ , or the probabilities. We can then write the generalized recovery equation set on the form

$$\Pi h^{-1}(\gamma) = [\delta \quad \delta^2 \quad \dots \quad \delta^T]^T, \tag{34}$$

where h is a one-to-one C^∞ smooth function of the parameter γ based on Eq. (33); see Appendix B for details.¹² Therefore, we are in the domain of Assumption 1* and, as long as $T > 2$ (since $N = 1$ is the number of risk aversion parameters and 2 is the total number of variables, δ and γ) then we know by Proposition 6 that the generalized recovery equation set generically has a unique solution.

We first seek to recover γ and δ by minimizing the pricing errors (again, see Appendix B for details). Panel A of Fig. 3 shows the objective function for this minimization problem. As seen from the figure, there is a unique solution to the problem, which naturally equals the true parameters $\hat{\delta} = 0.98, \hat{\gamma} = 4$.

Finally, we turn to the recovery of natural probabilities. It is worth noticing that we do *not* recover the Markov switching probabilities $\phi_{uu}, \phi_{dd}, \phi_{ud},$ or ϕ_{du} . Rather, what is recovered is the multi-period probabilities $p_{0,t}^{1,s_t}$ of transitioning from the initial state to each future state (consistent with the intuition conveyed in Fig. 1).¹³ The probabilities $p_{0,t}^{1,s_t}$ are recovered exactly. Fortunately, these multi-period probabilities are all we need for making predictions about such statistics as expected returns, variances, and quantiles across different time horizons.


7.2. Cox–Ross–Rubinstein and iid. consumption growth

We can capture the standard binomial model of Cox et al. (1979) (i.e., the discrete time counterpart to Black–Scholes–Merton) as follows. We consider the same model


¹¹ We note that prices of long-lived assets, for example the overall stock market, depends on both X_t and Y_t (even if the aggregate consumption Y_t is the aggregate dividend). Therefore, stock index options would provide information on Arrow–Debreu prices on each state $s_t = (y_t, x_t)$. Alternatively, we could consider recovery based only on Arrow–Debreu securities that depend on y_t . This would correspond to observing options on “dividend strips.” Either way, we get the same recovery results in the Mehra and Prescott (1985) model.

¹² Matlab code is available from the authors upon request.


¹³ Recovery of the underlying path-dependent probabilities is possible if we have access to Arrow–Debreu prices for all paths or if we assume that we know the structure of the underlying tree.



Panel A: Mehra Prescott (1985)



Panel B: iid. consumption



Panel C: Non-Markovian

Fig. 3. Generalized recovery: objective function in specific economic models. The figure shows the objective function used for the generalized recovery method, the squared pricing errors in (B.3). Panel A shows that the objective function for the Mehra and Prescott (1985) model has a unique minimum, making the generalized recovery feasible. Panel B shows that generalized recovery is not feasible in the Cox-Ross-Rubinstein model with iid. consumption, as the objective has a continuum of solutions. Panel C shows that generalized recovery is feasible in the non-Markovian model.

for aggregate consumption Y_t , but now X_t is iid. (corresponding to $\phi_{uu} = \phi_{du}$ and $\phi_{dd} = \phi_{ud}$). In other words, the standard binomial model has iid. consumption growth. Specifically, we assume that up and down probabilities are always 50% ($\phi_{uu} = \phi_{du} = \phi_{dd} = \phi_{ud} = 0.5$).

This binomial model implies a flat term structure that puts us in the case of Proposition 4, where recovery is impossible.¹⁴ Concretely, the problem is that the price matrix Π from Eq. (34) is not full rank. Hence, as seen in Fig. 3 Panel B, the objective of minimizing pricing errors has a continuum of solutions. In other words, recovery is not feasible.

7.3. A non-stationary model without Markov structure

Lastly, we consider a model where the consumption growth X_t is not Markov. Specifically, we still consider the binomial tree described above in Sections 7.1 and 7.2, but now we let the probability of transitioning up/down from any state s at any time t depend on the path taken from time 0 to time t . At each node at each path, we draw a random uniformly distributed probability for an up move, and, of course, assign one minus this probability to the next down node.

We now seek to recover δ and γ . As seen in Fig. 3 Panel C, the objective function has a unique solution that again equals the true parameters $\hat{\delta} = 0.98$ and $\hat{\gamma} = 4$. Hence, recovery can be possible even when the driving process is nonstationary and non-Markovian, again under parametric assumptions about the utility function (i.e., a model outside the scope of Ross 2015 and Borovicka et al. 2016).

8. Empirical analysis

This section describes our data, empirical methodology, and empirical findings.

8.1. Data and sample selection

We use the IvyDB database from OptionMetrics to extract information on standard call and put options written on the S&P 500 index for every last trading day of the month from January 1996 to December 2015. We obtain implied volatilities, strikes, and maturities, allowing us to back out market prices. As a proxy for the risk-free rate, we use the zero-coupon yield curve of the IvyDB database, which is derived from LIBOR rates and settlement prices of CME Eurodollar futures. We also obtain expected dividend payments, calculated under the assumption of a constant dividend yield over the lifetime of the option. We consider options with time to maturity between 10 and 360 days and apply standard filters, excluding contracts with zero open interest, zero trading volume, and quotes with best bid below \$0.50, and options with implied volatility higher than 100%.

¹⁴ iid. consumption growth and standard utility functions generally lead to a flat term structure because the price of a bond with τ periods to maturity can be written as $E_t(\delta^\tau \frac{u_{t+\tau}}{u_t}) = E_t(\prod_{s=1, \dots, \tau} \delta \frac{u_{t+s}}{u_{t+s-1}}) = (\frac{1}{1+r})^\tau$, where the expected utility increments are the same for all s because they depend on consumption growth $\frac{c_{t+s}}{c_{t+s-1}}$, which has constant expected value when it is iid.

8.2. Recovery methodology

The generalized recovery theorem relies on the knowledge of Arrow–Debreu state prices from the current initial state to all possible future states for several future time periods. Unfortunately, there is currently no market trading pure Arrow–Debreu securities. Therefore, we use options to back out Arrow–Debreu prices. Further, given the large number of states, we use the parametric kernel method from Section 6.

To study the robustness of recovery, we consider two different methods for backing out Arrow–Debreu prices and two different specifications of the pricing kernel, for a total of four different recovered distributions and preferences.

More specifically, we apply the following two methods of extracting Arrow–Debreu prices from options: (i) the parametric model of Bates (2000) and (ii) the non-parametric method of Jackwerth (2004). Each of the methods yields Arrow–Debreu prices across multiple time horizons and multiple index levels for each day t as described in detail in Appendix C.

Given these observed Arrow–Debreu prices, we recover preferences and probabilities based on the two different specifications of the pricing kernel that we denote “piecewise linear” and “polynomial” pricing kernels, respectively, as described in detail in Appendix D.

8.3. Computing statistics of the recovered distribution

Once we have recovered the probabilities of each state for each future time period, it is straightforward to compute any statistic under the physical probability distribution. If the level of the index at time t is S_t , then the state space consists of all integer values of the index between the minimum value $(1 - 2.5VIX_t)S_t$ and $(1 + 4VIX_t)S_t$. Let N_t denote the number of states as seen from time t , and think of state 1 as the lowest state and N_t as the highest state. We compute the recovered expected excess return μ_t at time t by summing over the N_t possible states:

$$\mu_t = E_t^{\mathbb{P}}[r_{t,t+1}] - r_{t,t+1}^f = \sum_{v=1}^{N_t} p_{t+1,v} r_{t+1,v} - r_{t,t+1}^f \tag{35}$$

where $r_{t,t+1}^f$ is the risk-free rate, $p_{t+1,v}$ is the recovered time- t conditional physical probability for the transition to state v at time $t + 1$, $r_{t+1,v} = \frac{S_{t+1}(v)}{S_t} - 1$ is the return in state v , and $S_{t+1}(v)$ is the value of the index at time $t + 1$ if state v is realized.

We compute the contemporaneous unpredictable innovation in the conditional expected return as

$$\Delta\mu_{t+1} = \mu_{t+1} - E_t[\mu_{t+1}], \tag{36}$$

where we impose an AR(1) process on the innovation to the risk premium $E_t[\mu_{t+1}] = \alpha_0 + \alpha_1\mu_t$ based on the regression

$$\mu_{t+1} = \alpha_0 + \alpha_1\mu_t + \varepsilon_{t+1}. \tag{37}$$

The estimated persistence parameter α_1 depends on the recovery method and ranges from 0.23 to 0.68 at the monthly horizon.

Table 1

Correlation matrix. This table shows the pairwise correlations between the recovered conditional expected excess return for different specifications of marginal utilities and method for estimating risk-neutral prices: (i) $\mu_{t,1}$: Bates and polynomial; (ii) $\mu_{t,2}$: Bates and piecewise linear; (iii) $\mu_{t,3}$: Jackwerth and polynomial; and (iv) $\mu_{t,4}$: Jackwerth and piecewise linear. We augment the table with pairwise correlations with the VIX_t index and the lower boundary on the equity premium, $SVIX_t$, due to Martin (2017).

	$\mu_{t,1}$	$\mu_{t,2}$	$\mu_{t,3}$	$\mu_{t,4}$	VIX_t	$SVIX_t$
$\mu_{t,1}$	1	0.31	0.31	0.54	0.39	0.39
$\mu_{t,2}$		1	0.23	0.38	0.41	0.38
$\mu_{t,3}$			1	0.37	0.34	0.31
$\mu_{t,4}$				1	0.32	0.29
VIX_t					1	0.93
$SVIX_t$						1

We compute the recovered conditional variance, $\text{VAR}_t^{\mathbb{P}}(r_{t,t+1})$, analogously to how we computed the expected return and we denote the recovered volatility by $\sigma_t = \sqrt{\text{VAR}_t^{\mathbb{P}}(r_{t,t+1})}$.

8.4. Empirical results

We next investigate the properties of the recovered probabilities based on each of our four methods. We first consider the recovered expected return. Table 1 shows the correlation matrix for the recovered expected returns based on each of our four methodologies as well as the VIX volatility index and the SVIX variable of Martin (2017). The good news is that all variables are positively correlated, as we would expect. The less good news is that the correlations between the different recovered expected returns are modest in magnitude, with an average pairwise correlation of only 0.35. This modest correlation is concerning because all these recovered expected returns should be measures of the same thing, namely the market’s expected return at any given time.

Fig. 4 shows the time series variation of the recovered expected return based on one of the methodologies (we plot just one time series since it is difficult to look at all four together). These recovered expected returns do not look unreasonable, but we next try to test their ability to predict actual realized returns. Specifically, we regress the ex post realized excess return on the ex ante recovered expected excess return, μ_t and the ex post innovation in expected return, $\Delta\mu_{t+1}$:

$$r_{t,t+1} = \beta_0 + \beta_1\mu_t + \beta_2\Delta\mu_{t+1} + \epsilon_{t,t+1}, \tag{38}$$

where $\epsilon_{t,t+1}$ is a noise term. To understand this regression, note that we are interested in testing whether the recovered probabilities give rise to reasonable expected returns, that is, time-varying risk premiums. For this, we want to test whether a higher ex ante expected return is associated with a higher ex post realized return ($\beta_1 > 0$), whether an increase in the risk premium is associated with a contemporaneous drop in the price ($\beta_2 < 0$), and whether the intercept as zero ($\beta_0 = 0$).

Table 2 reports the results of this regression for each of our four recovery methodologies as well as using VIX and SVIX as the expected return over the full sample from 1996




Fig. 4. Recovered conditional expected excess return. The figure plots monthly conditional expected excess market returns, recovered last trading day of each month from 1/1996 to 12/2015. Marginal utilities are polynomial in return and maturity and risk-neutral prices are estimated using Jackwerth (2004).

Table 2

Does the recovered expected return predict the future return? The table reports results of the regression of the ex post realized excess return r_{t+1} on the ex ante recovered expected excess return, μ_t , and the ex post innovation in expected return, $\Delta\mu_{t+1}$. In the last two columns we replace μ_t with the VIX or the SVIX of Martin (2017).

$$r_{t,t+1} = \beta_0 + \beta_1\mu_t + \beta_2\Delta\mu_{t+1} + \epsilon_{t,t+1}.$$

The regression uses monthly data over the full sample 1/1996–12/2015, t -statistics are reported in parentheses, and significance at a 10% level is indicated in bold.

Dependent variable	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$
Intercept	0.00 (0.05)	0.01 (2.69)	0.01 (2.34)	0.01 (2.64)	0.01 (2.05)	0.01 (1.78)
μ_t	0.63 (0.59)	-7.05 (-1.59)	-1.01 (-0.86)	-0.57 (-1.52)	-0.00 (-1.25)	0.18 (0.28)
$\Delta\mu_{t+1}$	-3.92 (-3.27)	-13.75 (-2.90)	-5.93 (-4.90)	-1.65 (-3.20)	-0.55 (-10.1)	-16.11 (-16.01)
Adj. R^2 (%)	3.6	3.6	8.8	4.3	30.0	51.7
Method:						
Expected excess return (μ_t)	Recovered	Recovered	Recovered	Recovered	VIX	SVIX
Q-prices	Bates	Bates	Jackwerth	Jackwerth		
Pricing kernel	Polynomial	Piecewise linear	Polynomial	Piecewise linear		

to 2015. First, the intercept β_0 is insignificantly different from zero using method 1 and using SVIX, but significantly different from zero in the other specifications, providing evidence against these models. Second, β_1 is insignificantly different from zero in all specifications, providing evidence against the models. The coefficient β_2 is highly significant and has the desired negative sign in all models. Further, as expected the absolute value of β_2 is greater than one since a shock to the discount rate leads to a larger shock to the price (cf. Gordon’s growth model for the extreme example of a permanent shock).

Table 3 reports the result of regression (38) over the subsample that excludes the global financial crisis (9/2008–7/2009), a subsample that has been considered in the literature (e.g., Martin 2017). The results here show little improvement. Method 1 is again the most encouraging in that β_0 is insignificantly different from zero, β_1 is positive albeit insignificant, and β_2 is significantly negative as expected. However, the other three methods only give the expected negative coefficient for β_2 . None of the estimates for β_1 are significant and all of the β_0 coefficients are significantly positive.

Finally, we consider the recovered physical volatility as plotted in Fig. 5. This recovered volatility looks reasonable. Further, the recovered volatilities are similar across the different methodologies with an average pairwise correlation of 0.98 and an average correlation to VIX of 0.92. It is not that surprising that volatilities can be recovered, but studying volatility provides a simple and powerful reality check of our method since the true future volatility is known with much less error than the expected return. Hence, we regress the ex post realized volatility on the ex ante recovered conditional volatility, σ_t :

$$\sqrt{\text{VAR}(r_{t,t+1})} = \beta_0 + \beta_1\sigma_t + \epsilon_{t,t+1}, \tag{39}$$

where the realized volatility $\sqrt{\text{VAR}(r_{t,t+1})}$ is computed using close-to-close daily data over the four weeks from t to $t + 1$ by OptionMetrics. We also run the same regression where we replace the recovered volatilities by the VIX volatility index. The theory predicts that $\beta_0 = 0$ and $\beta_1 = 1$.

Table 4 reports the results. As seen in Table 4, the estimated intercept coefficient β_0 is insignificant for models 1, 3, and 4, but it is significant for model 2. However, for

Table 3

Does the recovered expected return predict the future return – excluding 8/2008–7/2009. The table reports results of the regression of the ex post realized excess return r_{t+1} on the ex ante recovered expected excess return, μ_t , and the ex post innovation in expected return, $\Delta\mu_{t+1}$. In the last two columns we replace μ_t with the VIX or the SVIX of [Martin \(2017\)](#).

$$r_{t,t+1} = \beta_0 + \beta_1 \mu_t + \beta_2 \Delta\mu_{t+1} + \epsilon_{t,t+1}.$$

The regression uses monthly data over the full sample 1/1996–12/2015, t -statistics are reported in parentheses, and significance at a 10% level is indicated in bold.

Dependent variable	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$	$r_{t,t+1}$
Intercept	0.00 (0.30)	0.01 (2.56)	0.01 (1.88)	0.01 (2.52)	0.00 (1.11)	0.00 (0.12)
μ_t	1.46 (1.42)	-2.80 (-0.64)	1.06 (0.82)	-0.17 (-0.47)	0.00 (-0.25)	1.71 (1.99)
$\Delta\mu_{t+1}$	-4.38 (-3.84)	-12.74 (-2.77)	-6.23 (-5.19)	-2.11 (-4.37)	-0.50 (-8.75)	-17.69 (-15.53)
Adj. R^2 (%)	5.9	2.6	9.8	7.1	24.6	52.5
Method:						
Expected excess return (μ_t)	Recovered	Recovered	Recovered	Recovered	VIX	SVIX
Q-prices	Bates	Bates	Jackwerth	Jackwerth		
Pricing kernel	Polynomial	Piecewise linear	Polynomial	Piecewise linear		




Fig. 5. Recovered conditional volatility of excess return. The figure plots monthly conditional market volatility, recovered last trading day of each month from 1/1996 to 12/2015. Marginal utilities are polynomial in return and maturity and risk-neutral prices are estimated using [Jackwerth \(2004\)](#).

Table 4

Does the recovered volatility predict the future volatility? This table reports results of a monthly regression of the ex post realized volatility on the ex ante recovered return volatility, σ_t . In the last column we replace σ_t with the VIX.

$$\sqrt{\text{var}(r_{t,t+1})} = \beta_0 + \beta_1 \sigma_t + \epsilon_{t,t+1}.$$

The regression uses monthly data over the full sample 1/1996–12/2015, t -statistics are reported in parentheses, and significance at a 10% level is indicated in bold.

Dependent variable	$\sqrt{\text{var}(r_{t,t+1})}$	$\sqrt{\text{var}(r_{t,t+1})}$	$\sqrt{\text{var}(r_{t,t+1})}$	$\sqrt{\text{var}(r_{t,t+1})}$	$\sqrt{\text{var}(r_{t,t+1})}$
Intercept	-0.00 (-1.24)	-0.01 (-1.95)	-0.01 (-1.63)	-0.01 (-1.46)	-0.05 (-9.63)
σ_t	0.88 (16.67)	0.87 (16.89)	0.76 (16.14)	0.77 (15.78)	0.71 (17.19)
Adj. R^2 (%)	53.9	54.6	52.3	51.2	55.3
Method:					
Volatility (σ_t)	Recovered	Recovered	Recovered	Recovered	VIX
Q-prices	Bates	Bates	Jackwerth	Jackwerth	
Pricing kernel	Polynomial	Piecewise linear	Polynomial	Piecewise linear	

all models, the intercept is smaller than that of VIX, suggesting that the recovered volatilities are less biased than VIX.

The estimated slope coefficient β_1 is positive and highly significant for all models. Further, the estimated slope is close to the predicted value of 1, in particular closer than the estimated value for VIX. Lastly, we see that VIX has a slightly higher R^2 , which may reflect that the recovery method introduces some noise in the volatility measures.

In summary, we find substantial differences across the recovered probabilities based on different methodologies, and the predictive power for future returns appears weak both in the full sample and in the sample that excludes the global financial crisis. The recovered volatilities predict well the future volatility in a way that is less biased than VIX but slightly lower R^2 . We can reject that the recovered probabilities provide a perfect description of the future evolution of the market based on a Berkowitz (2001) test.¹⁵ This rejection could be due to the details of our implementation. For instance, while the true pricing kernel may depend on multiple factors, we assume that the state space is given by the level of S&P 500 since we do not observe option prices depending simultaneously on multiple factors.

9. Conclusion

We characterize when preferences and natural probabilities can be recovered from observed prices using a simple counting argument. We make no assumptions on the physical probability distribution, thus generalizing Ross (2015) who relies on strong time-homogeneity assumptions.

In economies with growth, our counting argument immediately shows that recovery is generally not feasible. While this finding parallels results by Borovicka et al. (2016), our intuitive counting argument is fundamentally different and does not rely on the assumptions of an infinite-period time-homogeneous Markov setting but rather is based on the general methods pioneered by Debreu (1970) for general equilibrium.

To pursue recovery even in economies with growth, e.g., classical multinomial models, we show how our method can be used when the pricing kernel can be parameterized by a sufficiently low-dimensional parameter vector. When recovery is feasible, our model allows a closed-form linearized solution. We implement our model empirically using several different specifications, testing the predictive power of the recovered statistics.

Acknowledgment

We are grateful for helpful comments from Jaroslav Borovicka, Peter Christoffersen, Horatio Cuesdeanu, Darrell

Duffie, Lars Peter Hansen, Jens Jackwerth, Pia Mølgaard, Jose Scheinkman, Paul Schneider, Paul Whelan as well as from seminar participants at London Business School, University of Zurich, University of Konstanz, University of Toronto, ATP, HEC-McGill winter finance workshop 2015, SED summer conference 2016, EFA 2016, AFA 2017. We are especially grateful to Stephen Ross (our AFA discussant) for several stimulating discussions, including encouraging us to study uniqueness, leading to new results in Proposition 2, and to Dan Petersen for guiding us through Sard's theorem. All three authors gratefully acknowledge support from the FRIC Center for Financial Frictions (grant no. DNR102) and Jensen and Pedersen from the European Research Council (ERC grant no. 312417).

Appendix A. Proofs

Proof of Proposition 1. We have already provided a proof for the statements 1 and 2 in the body of the text. Turning to statement 3, we note that the set X of all (δ, h, P) is a manifold-with-boundary of dimension $S \cdot T - T + S$. The discount rate, probabilities, and marginal utilities map into prices, which we denote by $F(\delta, h, P) = DPH = \Pi$, where, as before, $D = \text{diag}(\delta, \dots, \delta^T)$ and $H = \text{diag}(1, h_2, \dots, h_S)$, and F is C^∞ . If $S < T$, the image $F(X)$ has Lebesgue measure zero in $\mathbb{R}^{T \times S}$ by Sard's theorem, proving 3. Indeed, this means that the prices that are generated by the model $F(X)$ have measure zero relative to all prices Π .

Turning to statement 4, we first note that P and H can be uniquely recovered from (δ, Π) (given that Π is generically full rank). Indeed, H is recovered from Eq. (17) and P is recovered from Eq. (12). Therefore, we can focus on (δ, Π) .

For two different choices of the discount rate (δ_a, δ_b) and a single set of prices Π , we consider the triplet $(\delta_a, \delta_b, \Pi)$. We are interested in showing that the different discount rates cannot both be consistent with the same prices, generically. To show this, we consider the space M where the reverse is true, hoping to show that M is "small." Specifically, M is the set of triplets where Π is of full rank and both discount rates are consistent with the prices, that is, there exists (unique) P_i and H_i ($i = a, b$) such that $D_a P_a H_a = D_b P_b H_b = \Pi$.

Given that probabilities and marginal utilities can be uniquely recovered from prices and a discount rate (as explained above), we have a smooth map G from M to X by mapping any triplet $(\delta_a, \delta_b, \Pi)$ to (δ_a, h_a, P_a) , where (h_a, P_a) are the recovered marginal utility and probabilities. The image of this map consists exactly of those elements of X for which F is not injective. The proof is complete if we can show that this image has Lebesgue measure zero, which follows again by Sard's theorem if we can show that the dimension of M is strictly smaller than $ST - T + S$.

To study the dimension of M , we note that we can think of M as the space of triplets such that the span of Π contains both the points $(\delta_a, \delta_a^2, \dots, \delta_a^T)'$ and $(\delta_b, \delta_b^2, \dots, \delta_b^T)'$. The span of Π is given by $V_\Pi := \{\Pi \cdot (1, h_2, h_3, \dots, h_S)' | h_s > 0\}$, which is an affine $(S - 1)$ -dimensional subspace of \mathbb{R}^T for Π of full rank. The set of all those $\Pi \in \mathbb{R}^{T \times S}$ such that V_Π passes through two given points of \mathbb{R}^T (in general position with respect to

¹⁵ The details of this test are not reported for brevity. The idea is that given the estimated distribution \hat{f}_t of the excess return r_{t+1} at time t , the distribution of the transformed variable $u_{t+1} = \hat{f}_t(r_{t+1})$ should be uniform and the distribution of the further transformed variable $x_{t+1} = \Phi^{-1}(u_{t+1})$ should be standard normal, which is tested by estimated the coefficients in the model $x_{t+1} = c + \beta x_t + \epsilon_t$ and perform a likelihood ratio test of the joint hypothesis that $c = \beta = 0$ and $\text{Var}(\epsilon_t) = 1$.

each other) form a subspace of dimension $ST - 2(T - S + 1)$ since each point imposes $T - S + 1$ equations (and saying that the points are in general position means that all these equations are independent). Therefore, M is a manifold of dimension $ST - 2T + 2S$ since the pair (δ_a, δ_b) depends on two parameters, and for a given pair, there is a $(ST - 2T + 2S - 2)$ -dimensional subspace of possible Π (any two distinct points are always in general position). Hence, we see that $\dim(M) = ST - 2T + 2S < ST - T + S = \dim(X)$ since $S < T$, which implies that $G(M)$ has measure zero in X . Further, the prices where recovery is impossible, $F(G(M))$, have measure zero in the space of all prices generated by the model $F(X)$ where we use the Lebesgue measure on X to define a measure¹⁶ on $F(X)$. \square

Proof of Proposition 2. Let $\bar{\Pi}$ be an $S \times S$ transition matrix corresponding to an irreducible matrix (as in Ross). Without loss of generality we assume that the current state is the first state. Since prices are generated by a Ross economy, the observed matrix Π of multiperiod prices is given as

$$\Pi := \begin{pmatrix} (\bar{\Pi})_1 \\ (\bar{\Pi}^2)_1 \\ \vdots \\ (\bar{\Pi}^S)_1 \end{pmatrix}.$$

where $(\bar{\Pi})_1$ denotes the first row of $\bar{\Pi}$, $(\bar{\Pi}^2)_1$ is the first row of $\bar{\Pi}^2$, etc. We want to show that all solutions to the eigenvalue problem for $\bar{\Pi}$ give rise to solutions to our system (both the “correct solution” and the ones that, by the Perron–Frobenius theorem, do not generate viable solutions).

Observe that if $z = (z_1, \dots, z_S)'$ is a (right) eigenvector of $\bar{\Pi}$ with corresponding eigenvalue δ , then

$$\Pi z = (\delta z_1, \delta^2 z_2, \dots, \delta^S z_S)'.$$

If z is the eigenvector corresponding to the maximal eigenvalue of $\bar{\Pi}$, then we know that it is strictly positive. Generically, in the space of matrices $\bar{\Pi}$, the matrix is diagonalizable with eigenvectors that contain no zeros and with distinct nonzero eigenvalues—in particular, it has full rank. Therefore, generically, even for the other eigenvectors, we have that the coordinates of z are non-zero, so we can normalize z to have first coordinate 1. Now let the Ross probability matrix be defined (as in Ross)

$$\bar{P} = \frac{1}{\delta} \text{Diag}^{-1}(z) \bar{\Pi} \text{Diag}(z), \tag{A.1}$$

with corresponding multi-period probabilities given by

$$P := \begin{pmatrix} (\bar{P})_1 \\ (\bar{P}^2)_1 \\ \vdots \\ (\bar{P}^S)_1 \end{pmatrix}.$$

Note that since the rows of \bar{P} sum to 1, so do rows of P . Further, using (A.1),

$$P = \begin{pmatrix} \left(\frac{1}{\delta^1} \text{Diag}(z)^{-1} \bar{\Pi}^1 \text{Diag}(z)\right)_1 \\ \vdots \\ \left(\frac{1}{\delta^S} \text{Diag}(z)^{-1} \bar{\Pi}^S \text{Diag}(z)\right)_1 \end{pmatrix} = \begin{pmatrix} \left(\frac{1}{\delta^1} \bar{\Pi}^1 \text{Diag}(z)\right)_1 \\ \vdots \\ \left(\frac{1}{\delta^S} \bar{\Pi}^S \text{Diag}(z)\right)_1 \end{pmatrix} = D^{-1} \Pi \text{Diag}(z),$$

where the second equality uses that $z_1 = 1$, and that we only consider the first rows, and the last equation uses our maintained notation $D = \text{Diag}(\delta, \dots, \delta^S)$. We note that this equation is the same as our Eq. (12), which means that all solutions to Ross’s eigenvalue problem for the matrix $\bar{\Pi}$ also appear as solutions to our equations. The fact that P generated from the Ross solution \bar{P} is a solution to the generalized problem required no assumptions other than irreducibility, and this proves part 1 of the theorem.

To also obtain uniqueness of our solution, note that generically, there are S eigenvectors for Ross’s matrix from which a matrix P can be generated using (A.1). Each of these solutions can be used to generate a solution P to our problem, as shown above. The $S - 1$ solutions are “fake” in the sense that they imply that some marginal utilities (elements in the eigenvector z above) are negative. Hence, these solutions are also fake in the context of the generalized recovery framework. Given that Ross’s equations yield a total of S possible solutions to our problem, of which $S - 1$ are fake, we have a unique viable solution (by Proposition 1) if we can ensure that Π_{12} has full rank.

This follows from the generic property of $\bar{\Pi}$ as being diagonalizable with distinct, nonzero eigenvalues. In fact, we can show the stronger statement that Π has full rank: consider the diagonalization of Ross’s price matrix as $\bar{\Pi} = VZV'$, where $Z = \text{diag}(z_1, \dots, z_S)$ is the matrix of eigenvalues, and V is the matrix of eigenvectors. The k th row in the generalized-recovery pricing matrix is the first row (still assuming that the starting state is 1) of $\bar{\Pi}^k = VZ^kV'$. Letting v denote the first row in V , we see that the k th row of Π is $vZ^kV' = (v_1z_1^k, \dots, v_Sz_S^k)V'$ so

$$\Pi = \begin{bmatrix} 1 & \dots & 1 \\ \vdots & & \vdots \\ z_1^{T-1} & \dots & z_S^{T-1} \end{bmatrix} \begin{bmatrix} v_1z_1 & & 0 \\ & \ddots & \\ 0 & & v_Sz_S \end{bmatrix} V'. \tag{A.2}$$

Therefore, Π is full rank generically because it is the product of three full-rank matrices. Indeed, the first matrix is a Vandermonde matrix, which is full rank when the z ’s are nonzero and different, which is true generically. The second matrix is clearly also full rank since the v ’s are also nonzero generically, and the third matrix is full rank by construction. Hence our set of equations can have no more than S solutions, and since $S - 1$ of these are “fake”,

¹⁶ We can define a measure on $F(X)$ by $\mu^*(A) := \mu(F^{-1}(A))$ for any set A , where μ is the Lebesgue measure on X .

we have unique recovery of the solution corresponding to Ross's solution also, generically.

To see how to derive $\bar{\Pi}$ in an economy where Π arises from a time-homogeneous Ross economy, note that the following equation set must hold:

$$\underbrace{\begin{bmatrix} (\Pi)_2 \\ \vdots \\ (\Pi)_S \end{bmatrix}}_{(S-1) \times S} = \underbrace{\begin{bmatrix} (\Pi)_1 \\ \vdots \\ (\Pi)_{S-1} \end{bmatrix}}_{(S-1) \times S} \bar{\Pi}, \tag{A.3}$$

where $(\Pi)_i$ is the i th row of Π . Further, using the notation from Eq. (14) for blocks of Π and denoting the first row of $\bar{\Pi}$ by $\bar{\Pi}_1$ and remaining rows by $\bar{\Pi}_2$, we can rewrite this equation as

$$\begin{bmatrix} (\Pi)_2 \\ \vdots \\ (\Pi)_S \end{bmatrix} = \begin{bmatrix} \Pi_{11} & \Pi_{12} \end{bmatrix} \begin{bmatrix} \bar{\Pi}_1 \\ \bar{\Pi}_2 \end{bmatrix}. \tag{A.4}$$

Given that $\bar{\Pi}_1$ is known (because the one-period state prices from state 1 are observed), it is useful to further rewrite this system as

$$\begin{bmatrix} (\Pi)_2 \\ \vdots \\ (\Pi)_S \end{bmatrix} - \Pi_{11} \bar{\Pi}_1 = \Pi_{12} \bar{\Pi}_2. \tag{A.5}$$

Hence, when Π_{12} is full rank, the Ross price matrix $\bar{\Pi}_2$ can be derived uniquely and explicitly by premultiplying by $(\Pi_{12})^{-1}$. We have already shown in part 2, that Π_{12} has full rank generically. If Π_{12} does not have full rank, there exists a nonzero vector $v \in R^{S-1}$ for which $\Pi_{12}v = 0$. In this case, if we start from a solution for which $\bar{\Pi}_2$ has strictly positive elements, we can pick $\epsilon > 0$ small enough that adding ϵv to a row of $\bar{\Pi}_2$ yields a perturbed matrix $\bar{\Pi}_2^\epsilon$ whose elements are also strictly positive. Clearly, $\bar{\Pi}_2^\epsilon$ also satisfies (A.5), and hence the Ross price matrix is not unique, showing part 3. \square

Proof of Proposition 3. Consider first the case $S < T$. The dimension of the parameter set (transition probabilities + utility parameters) generating the generalized-recovery price matrix Π is $ST - T + S$, which is strictly greater than the dimension S^2 of the parameter space generating price matrices in Ross's homogeneous case. Hence, generically no time-homogeneous solution can generate a generalized recovery price Π .

Our framework is also more general in the case $S = T$. Recalling that $p_{\tau i}$ denotes the probability of going from the current state 1 to state i in τ periods, it is clear that in a time-homogeneous setting we must have $p_{22} \geq p_{11}p_{12}$, i.e., the probability of going from state 1 to state 2 in two periods is (conservatively) bounded below by the probability obtained by considering the particular path that stays in state 1 in the first time period and then jumps to state 2 in the second. However, such a bound need not apply for the true probabilities if the transition probabilities are not time-homogeneous. The set of parameters that can generate Π matrices that are not attainable from homogeneous transition probabilities is clearly of Lebesgue measure greater than zero in the S^2 -dimensional parameter space. \square

Proof of Proposition 4. Let R denote the diagonal matrix whose k th diagonal element is $\frac{1}{(1+r)^k}$. Having a flat term structure means that the matrix Π of state prices as seen from a particular starting state can be written as

$$\Pi = RQ,$$

which defines Q as a stochastic matrix (i.e., with rows that sum to 1). Clearly, by letting $\delta = 1/(1+r)$ and having risk neutrality, i.e. $H = I_S$ (the identity matrix of dimension S), we obtain a solution to our recovery problem

$$\Pi = RQ = DPH = RPI_S = RP,$$

by setting $P = Q$. \square

Proof of Proposition 5. The result follows from the following lemma. \square

Lemma 1. Suppose that $x^* \in \mathbb{R}^n$ is defined by $f(x^*) = 0$ for a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with full rank of the Jacobian df in the neighborhood of x^* , and x is defined as the solution to the equation, $f(\bar{x}) + df(\bar{x})(x - \bar{x}) = 0$, where f has been linearized around $\bar{x} = x^* + \Delta x \epsilon$ for $\Delta x \in \mathbb{R}^n$ and $\epsilon \in \mathbb{R}$. Then $x = x^* + o(\epsilon)$ for $\epsilon \rightarrow 0$.

Proof of Lemma 1. Since we have $x = \bar{x} - df^{-1}f(\bar{x})$ we see that, as $\epsilon \rightarrow 0$,

$$\frac{x - x^*}{\epsilon} = \frac{\bar{x} - x^*}{\epsilon} - df^{-1} \frac{f(\bar{x}) - f(x^*)}{\epsilon} \rightarrow \Delta x - df^{-1} df \Delta x = 0. \tag{A.6}$$

\square

Proof of Proposition 6. Following the same logic as the proof of Proposition 1, we note that the set X of all (δ, θ, P) is a manifold-with-boundary of dimension $S \cdot T - T + N + 1$. The discount rate, marginal utility parameters, and probabilities map into prices, which we denote by $F(\delta, \theta, P) = DPH = \Pi$, where as before, $D = \text{diag}(\delta, \dots, \delta^T)$ and $H = \text{diag}(h_1(\theta), h_2(\theta), \dots, h_S(\theta))$, and F is C^∞ . Since $N + 1 < T$, the image $F(X)$ has Lebesgue measure zero in $\mathbb{R}^{T \times S}$ by Sard's theorem, proving part 1.

Turning to part 2, we first note that P can be uniquely recovered from (θ, Π) using Eq. (12), where $\hat{\theta} = (\delta, \theta)$. Therefore, we can focus on (θ, Π) , studying the solutions to $\Pi(h_1^{-1}(\theta), \dots, h_S^{-1}(\theta))' = (\delta, \dots, \delta^T)'$.

For two different choices of the parameters $(\bar{\theta}_a, \bar{\theta}_b)$ and a single set of prices Π , we consider the triplet $(\bar{\theta}_a, \bar{\theta}_b, \Pi)$. We are interested in showing that the different parameters cannot both be consistent with the same prices, generically. To show this, we consider the space M where the reverse is true, hoping to show that M is "small." Specifically, M is the set of triplets where Π is of full rank and both discount rates are consistent with the prices, that is, there exists (unique) P_i ($i = a, b$) such that $D_a P_a H_a = D_b P_b H_b = \Pi$.

Given that probabilities can be uniquely recovered from prices and parameters, we have a smooth map G from M to X by mapping any triplet $(\bar{\theta}_a, \bar{\theta}_b, \Pi)$ to $(\delta_a, \theta_a, P_a)$. The image of this map consists exactly of those elements of X for which F is not injective. The proof is complete if we can show that this image has Lebesgue measure zero, which again follows by Sard's theorem if we can show that the dimension of M is strictly smaller than $S \cdot T - T + N + 1$.

To study the dimension of M , first consider $V_\Pi := \{\Pi(h_1^{-1}(\theta), \dots, h_S^{-1}(\theta))' | \theta \in \Theta\}$, which is an

N -dimensional submanifold of \mathbb{R}^T for Π of full rank and given that h is a one-to-one embedding. We note that we can think of M as the space of triplets such that V_Π contains both the points $(\delta_a, \delta_a^2, \dots, \delta_a^T)'$ and $(\delta_b, \delta_b^2, \dots, \delta_b^T)'$, where the corresponding θ 's are given uniquely from the definition of V_Π since Π is full rank and h is one-to-one. The set of all those $\Pi \in \mathbb{R}^{T \times S}$ such that V_Π passes through two given points of \mathbb{R}^T form a subspace of dimension $ST - 2(T - N)$ since each point imposes $T - N$ equations. Therefore, M is a manifold of dimension $ST - 2T + 2N + 2$. Hence, we see that $G(X)$ has measure zero in X and $F(G(X))$ has measure zero in $F(X)$. \square

Appendix B. Details on recovery in Mehra–Prescott

Let

$$\Pi = \begin{bmatrix} \pi_{0,1}^{0,d} & \pi_{0,1}^{1,u} & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & \pi_{0,2}^{0,d} & \pi_{0,2}^{1,d} & \pi_{0,2}^{1,u} & \pi_{0,2}^{2,u} & 0 & \dots & 0 & 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & \pi_{0,T}^{0,d} & \pi_{0,T}^{1,d} & \pi_{0,T}^{1,u} & \dots & \pi_{0,T}^{T,u} \end{bmatrix}, \tag{B.1}$$

where $\pi_{0,t}^{k,u}$ is the state price of making a total of k “up” moves in t periods where the last move was up, that is, the Arrow–Debreu price for the state $s_t = (y_t, x_t) = (u^k d^{t-k}, u)$. Similarly, $\pi_{0,t}^{k,d}$ is the state price of making a total of k up moves in t periods where the last move was down.

Π has dimension $T \times (\sum_{t=1}^T 2t)$. This implies that the $h^{-1}(\gamma)$ vector of inverse marginal utility ratios must be $(\sum_{t=1}^T 2t)$ -dimensional. We fix this in the following way. We let

$$h^{-1}(\gamma) = \begin{bmatrix} (y_1^0)^\gamma & (y_1^1)^\gamma & (y_2^0)^\gamma & (y_2^1)^\gamma & (y_2^1)^\gamma \\ (y_2^2)^\gamma & \dots & (y_T^T)^\gamma \end{bmatrix}, \tag{B.2}$$

where $y_t^k = u^k d^{t-k}$ is the level of aggregate consumption when making a total of k up moves in t periods, and γ is the risk-aversion parameter that we wish to recover.

There is no closed-form solution to the nonlinear case of CRRA preferences. To obtain model estimates we sort to a numerical exercise, which is to minimize the objective function g :

$$\min_{\gamma, \delta} g(\gamma, \delta) := \text{norm} \left(\Pi h^{-1}(\gamma) - \begin{bmatrix} \delta \\ \delta^2 \\ \vdots \\ \delta^T \end{bmatrix} \right) \tag{B.3}$$

s.t. $\gamma \in \mathbb{R}_+$
 $\delta \in (0, 1]$

Based on the recovered (γ, δ) that solve this minimization problem, we can recover the natural probabilities from Eq. (33).

Appendix C. Computing state prices empirically

Before we can recover probabilities, we need to know the Arrow–Debreu prices or, said differently, characterize the risk-neutral distribution. There exist many ways to

do this in practice based on observed option prices, including various interpolation methods. We implement two methods: (i) the parametric stochastic volatility model of Bates (2000) and (ii) the non-parametric “fast and stable” method of Jackwerth (2004).¹⁷

C.1. The Bates (2000) stochastic volatility model with jumps

To ensure that we start with an arbitrage-free collection of Arrow–Debreu prices by strike and maturity, we use the model of Bates (2000) to derive state prices from observed option prices. This parametric approach puts structure on the tails of the risk-neutral density, which also allows us to extrapolate outside the range of observable option quotes. While the Bates (2000) model may not be the true

specification of the economy, we simply use this framework as a standard method in the literature to compute state prices, and, consistent with this pragmatic view, we allow parameters to change over time (which also avoids look-ahead bias).

In this model, the risk-neutral process for the price of the underlying asset, S_t , and the instantaneous variance, V_t , are assumed to be of the form

$$dS_t/S_t = (r^f - d - \lambda \bar{k})dt + \sqrt{V_t}dZ_t + kdq_t \tag{C.1}$$

$$dV_t = (\alpha - \beta V_t)dt + \sigma_v \sqrt{V_t}dZ_{vt}, \tag{C.2}$$

where Z_t and Z_{vt} are Brownian motions with correlation ρ , and q_t is a Poisson counting process that captures the risk of jumps in the price. The jumps occur with intensity λ , and each jump causes the price to be multiplied by the factor $1 + k$, which is lognormally distributed, i.e., $\ln(1 + k) \sim N(\ln(1 + \bar{k}) + \frac{1}{2}\delta^2, \delta^2)$. Further, r^f is the risk-free rate, and d is the dividend yield.

We calibrate these model parameters every fourth Wednesday as follows.¹⁸ On each day, given the current level of the market S_t and the risk-free term structure $r_{t,t+\tau}^f$, we find the model parameters $(\alpha, \beta, \lambda, \bar{k}, \sigma_v, \delta)$ and state variable V_t that minimize the vega-weighted squared pricing errors for 50 call and put options, following the methodology of Trolle and Schwartz (2009). The 50 chosen

¹⁷ Our estimated state prices are reported on the corresponding author’s personal website. Arriving at these estimates requires some ad-hoc adjustments, such as elimination of outliers. We do not provide a full description of these adjustments as our focus is on the transition from state prices to recovered probabilities.

¹⁸ We use data for every fourth Wednesday as a compromise between (i) the tradition in the asset pricing literature on return predictability of focusing on monthly returns and (ii) the tradition in the option literature of focusing on Wednesdays, where among other reasons option liquidity is high.

call/put options are those with the highest volumes. We allow the model parameters to vary over time since we simply use the model to smooth observed option prices (that may be noisy) such that they are arbitrage-free.

Once we have obtained model estimates, we compute the risk-neutral density $f(\tau, S_\tau)$ for any time τ periods into the future and state S_τ given the current time state S_t as:

$$f(\tau; S_\tau) = \frac{1}{\pi} \int_0^\infty \left(\frac{S_\tau}{S_t}\right)^{-iu} \psi(\tau, u) du, \tag{C.3}$$

that is, by integrating the characteristic function ψ numerically using the Gauss–Laguerre quadrature method. Knowing the risk-neutral density, the corresponding state price density $\pi(\tau; S_\tau)$ is the density discounted by the τ -period risk-free rate $r_{t,t+\tau}^f$:

$$\pi(\tau; S_\tau) = e^{-r_{t,t+\tau}^f} f(\tau; S_\tau). \tag{C.4}$$

This completes the computation of state prices. Indeed, we think of $\pi(\tau; S_\tau)$ as the Arrow–Debreu prices we need as a starting point for our recovery for each index level. For example, $\pi(1, 2000)$ is the Arrow–Debreu price of receiving \$1 in one year if the S&P 500 is between 2000 and 2001. We consider the grid of maturities and index levels described in Section 8.2.

C.2. The Jackwerth (2004) fast and stable method

We are interested in converting a (noisy) sparse set of implied volatilities into a full risk-neutral distribution. In Section C.1 we imposed a parametric form on the implied volatility surface through a stochastic volatility model with jumps. In this section we refrain from imposing any structure on implied volatilities, that is, we fit a nonparametric method to implied volatilities. The method we have chosen is the fast and stable method of Jackwerth (2004). This method has a single tuning parameter, λ , which simultaneously controls the smoothness of the function and the fit to observed implied volatilities. Clearly, there is a trade-off in choosing the value of the tuning parameter, which is the smoother the function, the worse the fit to observations. We therefore control the smoothness of the fit by imposing two conditions: (i) the estimated implied volatilities gives rise to a nonnegative risk-neutral distribution, and (ii) the risk-neutral distribution is unimodal in the range from 0.8 to 1.2 in moneyness (defined as S_t/S_0 , the index level at time t relative to the current index level). Under these conditions we minimize the objective function:

$$\min_{\sigma_s} \frac{1}{2(S+1)} \sum_{s=1}^S (\sigma_s'')^2 + \frac{\lambda}{2I} \sum_{i=1}^I (\sigma_i - \bar{\sigma}_i)^2, \tag{C.5}$$

where S is the number of states. σ_s is the implied volatility associated with state s . σ_s'' is the second derivative of the implied volatility function with respect to strike prices. $i = 1, \dots, I$ is the index for the observed implied volatilities, and $\bar{\sigma}_i$ is the i th observed implied volatility. As seen from Eq. (C.5), if λ is high then the fit to observations will be good compared to when λ is low. We therefore choose the highest value of λ which satisfies our two conditions described above. See Jackwerth (2004) for further comments on the method.

Once a smooth function for the implied volatilities is obtained we can back out a risk-neutral distribution by evaluating the Black and Scholes (1973) formula in the estimated implied volatilities and then differentiate the resulting call function twice with respect to strike prices as explained in Breeden and Litzenberger (1978).

The fast and stable method estimates a single option maturity at a time. In the period from January 1996 until December 2015, we have at least seven maturities on any given last trading day of the month. In the framework of Proposition 6, this allows us to parameterize the pricing kernel with up to six parameters and still obtain generalized recovery.

Appendix D. Pricing kernels used in empirical analysis

Piecewise linear. The inverse marginal utilities are piecewise linear over states. Given the initial state 1 at time 0 the τ -period inverse marginal utility ratio in state s is:¹⁹

$$(h_s^\tau(\theta))^{-1} = B_s \theta. \tag{D.1}$$

Here θ is an N -dimensional column vector, and B_s is the s 'th row of the known $S \times N$ “design matrix” B . In our empirical implementation N is 5.²⁰ Interpreting the parameters $\theta_1, \dots, \theta_N$ we let the first parameter θ_1 determine the initial level of the inverse pricing kernel $H^{-1}e = B\theta$. The next parameter, θ_2 , determines the initial slope of the first line segment. Similarly, θ_3 is the slope of the next line segment generated by $B\theta$.

We impose that $\theta_1, \dots, \theta_N \geq 0$, which means that the inverse pricing kernel is monotonically increasing or, equivalently, that the pricing kernel is monotonically decreasing i.e., that marginal utility decreases at higher levels of wealth.

The design matrix is characterized by its break points that separate the state space into $N - 2$ regions. These regions are chosen as follows. The lowest region ranges over states from $(1 - 2.5SVIX_t)S_t$ to $(1 - SVIX_t)S_t$, where S_t is the current (time t) level of the S&P 500 index and $SVIX$ is the risk-neutral variance used by Martin (2017). The highest region covers states ranging from $(1 + SVIX_t)S_t$ to $(1 + 4VIX_t)S_t$. In between these extremes, we consider $N - 3$ regions of equal size in the range $(1 - SVIX_t)S_0$ to $(1 + SVIX_t)S_t$. When using this specification of B and the estimated Arrow–Debreu prices, we obtain an $S \times N$ matrix ΠB with full rank for every last trading day of the month for the period 1/1996 to 12/2015.

With this in place we set up the following minimization problem

$$\begin{aligned} \min_{\theta, \delta} \text{norm}(D^{-1} \Pi B \theta - 1) \\ \text{s.t. } \theta > 0 \\ \delta \in (0, 1] \end{aligned} \tag{D.2}$$

¹⁹ Notice again that $(h_s^\tau(\theta))^{-1} = \frac{1}{h_s^\tau(\theta)}$ and is not the inverse function.

²⁰ The lowest number of maturities with observed option prices in our sample is seven. Therefore, we can impose a structure on the pricing kernel with at most six parameters, and hence N can at most be five because of the sixth parameter δ .

Given a state price matrix Π and a design matrix B , we estimate the θ and δ that best fit the model in a squared error sense. Once the marginal utilities and discount rate have been recovered, we back out the multi-period physical probabilities as

$$P = D^{-1} \Pi \text{diag}(B\theta), \tag{D.3}$$

where D is a diagonal matrix with elements $D_{ii} = \delta^i$, and $\text{diag}(B\theta)$ is a diagonal matrix with elements $\text{diag}(B\theta)_{jj} = B_j\theta$, where B_j is the j th row of B . We normalize P to have row sums of one, which is necessary since θ and δ are found from the minimization problem in (D.2) and are not solved perfectly.

Polynomial. The inverse marginal utility ratio is a polynomial in the return on the market and time horizon. Given the initial state 1 at time 0 the τ -period inverse marginal utility ratio in state s is:

$$(h_s^\tau(\theta))^{-1} = \beta_0 + \beta_1 r_s + \beta_2 r_s^2 + \beta_3 \tau r_s + \beta_4 \tau r_s^2. \tag{D.4}$$

Here $r_s = S_s/S_1 - 1$ is the return on the market in state s . The parameters of interest are $\theta = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$. In our implementation we impose three conditions on the parameters: (i) $\beta_0 > 0$, ensuring a positive pricing kernel when $r = 0$; (ii) the risk-premium is non-negative, and (iii) the inverse marginal utility ratios are always strictly positive (we set a lower bound on the inverse marginal utility ratio at 0.01). This means that the parameters $\beta_1, \beta_2, \beta_3, \beta_4$ can move freely (within the space of the conditions) and are all allowed to be either positive or negative.

The polynomial specification of the inverse marginal utility ratios illustrates one possible way of imposing structure on the marginal utilities, not only in the state dimension but also in the time horizon dimension. This specification allows marginal utilities in a given state, say s , to differ when considering different time horizons, that is, e.g. $h_s^\tau(\theta) \neq h_s^{\tau+1}(\theta)$. The polynomial specification nests the linear specification as a special case when β_2, β_3 , and β_4 are all zero.

The minimization procedure for the polynomial specification is

$$\begin{aligned} \min_{\theta, \delta} & \sum_{t=1}^T \left(\left(\sum_{s=1}^S \delta^{-t} \pi_{ts} (h_s^\tau(\theta))^{-1} \right) - 1 \right)^2 \\ \text{s.t.} & \quad \beta_0 > 0 \\ & \quad E_0^P(r_t | \theta, \delta) - r_t^f \geq 0 \quad \text{for all } t \in (1, \dots, T) \\ & \quad (h_s^\tau(\theta))^{-1} > 0 \quad \text{for all } s \in (1, \dots, S) \\ & \quad \text{and all } \tau \in (1, \dots, T) \\ & \quad \delta \in (0, 1] \end{aligned} \tag{D.5}$$

where π_{ts} is the state price in state s with time horizon t . Here $E_0^P(r_t | \theta, \delta) - r_t^f$ is the excess return given parameter values $\theta = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$, and δ . We also impose that $\beta_1 > 0$ and $\beta_2 < 0$, these parameter restrictions help ensure a positive risk premium.

Given estimates of $\delta, \beta_0, \beta_1, \beta_2, \beta_3$, and β_4 , we can arrive at the t period physical probabilities as

$$P_t = \delta^{-t} \Pi_t \text{diag} \left((h_s^\tau(\theta))^{-1} \right), \tag{D.6}$$

where Π_t is the t th row of the state price matrix Π , and r is an $S \times 1$ -dimensional vector of returns over states. We normalize P to have row sums of one; this is necessary since θ and δ are found from the minimization problem in (D.5) and are not solved perfectly.

References

Ait-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. *J. Econom.* 94, 9–51.
 Audrino, F., Huitema, R., Ludwig, M., 2014. An Empirical Analysis of the Ross Recovery Theorem. Unpublished working paper.
 Bakshi, G., Chabi-Yo, F., Gao, X., 2015. A recovery that we can trust? Deducing and testing the restrictions of the recovery theorem. *Rev. Financ. Stud.* 31, 532–555.
 Bates, D.S., 2000. Post-'87 crash fears in the S&P 500 futures options market. *J. Econom.* 94, 181–238.
 Berkowitz, J., 2001. Testing density forecasts, with applications to risk management. *J. Bus. Econ. Stat.* 19, 465–474.
 Black, F., Scholes, M.S., 1973. The pricing of options and corporate liabilities. *J. Polit. Econ.* 81, 637–654.
 Bliss, R.R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. *J. Financ.* 1, 407–446.
 Bollerslev, T., Todorov, V., 2011. Tails, fears, and risk premia. *J. Financ.* 66, 2165–2211.
 Borovicka, J., Hansen, L.P., Scheinkman, J.A., 2016. Misspecified recovery. *J. Financ.* 71, 2493–2544.
 Breeden, D.T., Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. *J. Bus.* 51, 621–651.
 Carr, P., Yu, J., 2012. Risk, return, and Ross recovery. *J. Deriv.* 20, 38–59.
 Christoffersen, P., Heston, S.L., Jacobs, K., 2013. Capturing option anomalies with a variance-dependent pricing kernel. *Rev. Financ. Stud.* 26, 1963–2006.
 Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option pricing: a simplified approach. *J. Financ. Econ.* 7, 229–263.
 Debreu, G., 1970. Economies with a finite set of equilibria. *Econometrica* 38 (3), 387–392.
 Hansen, L.P., Scheinkman, J.A., 2009. Long-term risk: an operator approach. *Econometrica* 77, 177–234.
 Jackwerth, J., 2000. Recovering risk aversion from option prices and realized returns. *Rev. Financ. Stud.* 13, 433–451.
 Jackwerth, J.C., 2004. Option-Implied Risk-Neutral Distributions and Risk Aversion. CFA Institute Research Foundation of AIMR Publications, pp. 1–86.
 Jackwerth, J.C., Rubinstein, M., 1996. Recovering probability distributions from option prices. *J. Financ.* 51, 1611–1631.
 Kazemi, H., 1992. An intertemporal model of asset prices in a Markov economy with a limiting stationary distribution. *Rev. Financ. Stud.* 5, 85–104.
 Linetsky, V., Qin, L., 2016. Positive eigenfunctions of Markovian pricing operators: Hansen–Scheinkman factorization, Ross recovery, and long-term pricing. *Oper. Res.* 64, 99–117.
 Malamud, S., 2016. Noisy Recovery. Unpublished working paper. Swiss Finance Institute.
 Martin, I., 2017. What is the expected return on the market? *Q. J. Econ.* 132, 367–433.
 Martin, I., Ross, S., 2018. Notes on the yield curve. *J. Financ. Econ.* forthcoming.
 Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. *J. Monet. Econ.* 15, 145–161.
 Qin, L., Linetsky, V., 2017. Long-term risk: a martingale approach. *Econometrica* 85, 299–312.
 Rosenberg, J.V., Engle, R.F., 2002. Empirical pricing kernels. *J. Financ. Econ.* 64, 341–372.
 Ross, S., 2015. The recovery theorem. *J. Financ.* 70, 615–648.
 Schneider, P., Trojani, F., 2018. (Almost) Model-free recovery. *J. Financ.* forthcoming.
 Trolle, A.B., Schwartz, E.S., 2009. Unspanned stochastic volatility and the pricing of commodity derivatives. *Rev. Financ. Stud.* 22, 4423–4461.
 Walden, J., 2017. Recovery with unbounded diffusion processes. *Rev. Financ. Stud.* 21, 1403–1444.