Does the Ross Recovery Theorem work Empirically?

Jens Carsten Jackwerth * Marco Menner †

January 15, 2018

Abstract

Starting with the fundamental relationship that state prices are the product of physical probabilities and the pricing kernel, Ross (2015) shows that, given strong assumptions, knowing state prices suffices for backing out physical probabilities and the pricing kernel at the same time. We find that such recovered physical distributions based on the S&P 500 index are incompatible with future realized returns. This negative result remains even when we add economically reasonable constraints. Reasons for the rejection seem to be numerical instabilities of the recovery algorithm and the inability of the constrained versions to generate pricing kernels sufficiently away from risk-neutrality.

Keywords: Ross recovery, pricing kernel, risk-neutral density, transition state prices, physical probabilities

---

*Jens Jackwerth is from the University of Konstanz, PO Box 134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-2196, Fax: +49-(0)7531-88-3120, jens.jackwerth@uni-konstanz.de

†Marco Menner is from the University of Konstanz, PO Box 134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-3346, Fax: +49-(0)7531-88-3120, marco.menner@uni-konstanz.de

We received helpful comments and suggestions from Guenther Franke, Anisha Ghosh, Eric Renault, Christian Schlag, and Bjorn Eraker. We thank participants at the 2018 American Finance Association Meeting in Philadelphia, the ESSFM Workshop at Gerzensee and at the DGF Conference in Bonn as well as seminar participants at the University of Konstanz, the University of Strasbourg, the University of Zurich, the University of St. Gallen, the University of Sydney, and the University of Queensland.

I. Introduction

Much of financial economics revolves around the triangular relation between physical return probabilities p, which are state prices π divided by the pricing kernel m:[1]

$$\text{physical probability } p = \frac{\text{state price } \pi}{\text{pricing kernel } m} \tag{1}$$

Researchers typically pick any two variables to find the third. In option pricing, for example, physical probabilities are changed into risk-neutral ones, which are nothing but normalized state prices, via the pricing kernel. Differently, the pricing kernel puzzle literature, e.g. as surveyed in Cuesdeanu and Jackwerth (2016), starts out with risk-neutral and physical probabilities in order to find empirical pricing kernels. Yet Ross (2015) presents a recovery theorem, which allows to back out both the pricing kernel and physical probabilities by only using state prices. To achieve this amazing feat, he needs to make strong assumptions concerning the economy. We investigate his claim and test if the recovered physical probabilities are compatible with future realized S&P 500 returns. We further analyze if the shape of the recovered pricing kernel is in line with utility theory. To understand our sobering results, we discuss in detail why the recovery theorem does not perform well empirically.

The Ross (2015) recovery theorem is based on three assumptions. First, it requires time-homogeneous transition state prices $\pi_{i,j}$ that represent state prices of moving from any given state i today to any other state j in the future. Such transition state prices include the usual *spot* state prices $\pi_{0,j}$ with 0 representing the current state of the economy. Spot state prices can be readily found from option prices, see e.g. Jackwerth (2004). Yet Ross recovery also requires as inputs the transition state prices emanating from alternative, hypothetical states of the world.[2] There is, however, no market with actively traded financial products, that allows backing out the (non-spot) transition state prices $\pi_{i,j}$ directly. We will therefore use information on spot state prices with different maturities to obtain transition state prices

[1] See e.g. Cochrane (2000), pp. 50. A state (Arrow-Debreu) price represents the dollar amount an investor is willing to pay for a security that pays out one dollar if the particular state occurs and nothing if any other state occurs. The pricing kernel is closely related to the marginal utility of such investor.

[2] Imagine that the current state of the world is characterized by the S&P 500 being at 1000. Let there be two future states, 900 and 1000, to which the spot state prices (emanating from 1000) relate. The required other transition state prices are the ones emanating (hypothetically) from 900 and ending at 900 or 1000 one period later.

linking N states. This requires us to first estimate an interpolated spot state price surface from sparse option data. There is no unique way for estimating transition state prices $\pi_{i,j}$ from spot state prices, and so we suggest several different approaches. We also introduce a version that allows us to apply the recovery theorem directly to spot state prices $\pi_{0,j}$ without the need to estimate all the other transition state prices.

Second, all transition state prices need to be positive, which turns out to be a fairly benign assumption as we can always force some small positive state price.[3]

Finally, the pricing kernel is restricted to be a constant times the ratio of values in state j over values in state i. A convenient economic interpretation of these values is to associate them with marginal utilities in those states. The constant can then be associated with a utility discount factor.

Taken together, the three assumptions allow Ross (2015) to formulate a unique eigenvalue problem. Its solution yields the physical transition probabilities $p_{i,j}$, which represent physical probabilities of moving from state i to state j, and the pricing kernel.

After we achieve recovery, we empirically test the hypothesis that future realized S&P 500 returns are drawn from the recovered physical spot distribution $p_{0,j}$. At each date τ, we work out the percentile of next month's S&P 500 return based on the recovered physical spot cumulative distribution function. This leaves us with one value x_τ for each date τ that lies in between 0 and 1. If our hypothesis holds, the set $\{x_\tau\}$ is uniformly distributed. We can strongly reject our hypothesis for three different statistical tests: The Berkowitz (2001) test, the uniformity test introduced by Knüppel (2015), and the Kolmogorov-Smirnov test.[4] Thus, Ross (2015) recovers physical spot distributions of returns, which are inconsistent with future realized S&P 500 returns. We further find that the recovery theorem does not produce downward sloping pricing kernels (as one would expect based on risk averse preferences) but that they are riddled with local minima and maxima. In contrast, we cannot reject the hypotheses that future S&P 500 returns are drawn from simple physical distributions based on a power pricing kernel or the five-year historical return distribution.

The empirical problems of Ross recovery stem from several sources. First, it is hard to obtain transition state prices from option prices. If we use a basic implementation of Ross recovery, which requires no additional assumptions besides positivity of transition state prices,

[3]Technically, some zero values could be allowed as long as any state can still be reached from any other state, possibly via some other intermediate states, Ross (2015).

[4]Note that Berkowitz (2001) does not directly test for uniformity of the values but for standard normality of a standard normal transformation of the values.

we obtain unstable transition state prices that, in addition, exhibit unrealistic properties such as multi-modality and imply extremely high or low risk-free rates in different states. Yet, if we introduce economically reasonable constraints to calm down the transition state prices (starting with a mild constraint of bounded state-dependent risk-free rates and proceeding with an additional constraint of unimodal transition state prices) then the recovery theorem generates almost flat pricing kernels.

Second, we argue that the strong assumption concerning the functional form of the pricing kernel is rather limiting.[5] We further highlight that the recovered pricing kernel is highly dependent on the structure of the transition state price matrix, which in turn is not well identified from the option prices.

Third, the assumption of time-homogeneous transition state prices might not hold and different periods may require different transition state prices. Indeed, we show empirically that the simultaneous fit to short- and long-dated options is poor for Ross recovery.

Fourth, we relate the recovered pricing kernel to two theoretical pricing kernel components, which we take from the following literature. Pre-dating Ross (2015), Hansen and Scheinkman (2009) already showed that, given the underlying Markovian environment, the Perron-Frobenius theorem can be used to recover probabilities p. They further argue those probabilities may provide useful long-term insights into risk pricing.[6] One can then relate the probabilities p via a multiplicative adjustment to the spot state prices. We follow Alvarez and Jermann (2005) and call this adjustment the transitory pricing kernel component.

Further, Hansen and Scheinkman (2009) relate the true physical probabilities via another multiplicative adjustment to the probabilities p. We again follow Alvarez and Jermann (2005) and call this adjustment the permanent pricing kernel component.[7] The total pricing kernel is thus the product of the permanent and the transitory components.

Borovicka et al. (2016) show that Ross (2015) recovers the probabilities p and, by implicitly setting the permanent pricing kernel component to one, Ross interprets p as the true physical probabilities. Yet, is the permanent component truly one? Bakshi et al. (2016) use options on 30-year Treasury bond futures and solve convex minimization problems to extract

[5]Note that standard CARA or CRRA pricing kernels are incompatible with Ross recovery as his multi-period pricing kernels are identical but for a constant, while, under CARA or CRRA, they are multiplications of single-period pricing kernels.

[6]Sharing this long-term perspective, Martin and Ross (2013) show that the pricing kernel implied by Ross recovery depends on the long end of the yield curve.

[7]Other authors use the term martingale component.

the permanent pricing kernel component. They find that this component features considerable dispersion and is not a constant value of one, contrary to the implicit assumption made by Ross.[8]

We presently return to our fourth empirical problem of Ross recovery. The probabilities p from Ross recovery may thus differ from the true physical probabilities and our paper shows that they indeed do. As Ross recovers the probabilities p (and not the true physical probabilities), Ross recovery should allow us to extract the transitory pricing kernel component. Alternatively, Bakshi et al. (2016) extract the transitory component of the pricing kernel using data on 30-year Treasury bond futures. We can thus compare the two approaches. We obtain a time series of realizations of the (transitory) pricing kernel (component) from Ross recovery and we implement the approach of Bakshi et al. (2016). Relating the two approaches to each other, we show that they are not the same. We view this as further evidence that Ross recovery does not work well empirically.[9]

Overall, we add to the theoretical literature, which has been critical of Ross recovery, by pinpointing exactly where Ross recovery goes awry. Moreover, we answer the intriguing question if Ross recovery, despite all theoretical short-comings, might still be useful empirically as a rough approximation of reality. Alas, our work confirms the negative theoretical outlook.

Only a few papers already investigate the recovery theorem from an empirical perspective. Closest to our work are Audrino et al. (2015), who also implement Ross recovery on S&P 500 index options. Their recovered pricing kernels tend to be rather smooth and U-shaped, as opposed to our wavy pricing kernels. This surprising behaviour seems to be due to a particular modelling choice for a penalty term.[10] Their further empirical focus is on developing profitable trading strategies based on the recovered physical probabilities, but, unlike our work, they do not statistically test if future realized returns are drawn from the recovered distribution.

Also informing our empirical work is a simulation study by Tran and Xia (2014), who show that the recovered probabilities vary substantially for different state space dimensions.

[8]Christensen (2017) also provides a non-parametric empirical framework to estimate a time series of the permanent and the transitory pricing kernel components using both equity and macroeconomic data.

[9]We thank Anisha Ghosh for suggesting this study.

[10]The subtle reason is a quadratic penalty term, which they use to force all transition state prices to zero. This penalty is stronger for states further away from the current states as option prices are more sensitive to state prices around the current state. As a result, the implied risk-free rates (=1 - sum of transition state prices) increase in the distance to the current state, which in turn leads to U-shaped pricing kernels.

We check for this prediction in our robustness tests but it empirically does not matter much for our results.

As an alternative to the numerically difficult recovery of transition state prices from spot state prices, we suggest an additional, implicit method, which obviates that recovery and works directly with the spot state prices. In independent work, Jensen et al. (2017) suggest the same method and add further restrictions of the pricing kernel. They then focus on analyzing the theoretical properties of their generalized recovery. In a short empirical study, they use the mean of the recovered physical probabilities to predict S&P 500 return with a significant R^2 of 1.28%. They also apply a Berkowitz test and reject that future realized S&P 500 returns are drawn from the recovered distribution based on their particular model.

Further work on the implementation of Ross recovery is Massacci et al. (2016), whose fast non-linear programming approach allows for economic constraints such as positive state-dependent risk-free rates and the unimodality of transition state prices.

To our knowledge, we are the first to show empirically that several plausible implementations of the Ross recovery theorem are not compatible with future realized returns of the S&P 500 index and to analyze why this is the case.

While Ross recovery works on a discrete state space using the Perron-Frobenius theorem, Carr and Yu (2012) show that recovery can be achieved in continuous time for univariate time-homogeneous bounded diffusion process by using Sturm-Liouville theory. Walden (2016) further investigates an extension of the recovery theorem to continuous time if the diffusion process is unbounded. He derives necessary and sufficient conditions that enable recovery and finds that recovery is still possible for many of these unbounded processes. Additional works on Ross recovery in continuous time are Qin and Linetsky (2016), Qin et al. (2016), and Dubynskiy and Goldstein (2013). Related to the recovery literature, Schneider and Trojani (2016) extract physical moments from spot state prices in a unique minimum variance pricing kernel framework with mild economic assumptions on specific risk premia.

The remainder of the paper proceeds as follows. Section II explains the Ross recovery theorem. In Section III, we introduce our methods to obtain spot state prices, to back out transition state prices, and to apply the theorem without using transition state prices. We further explain the Berkowitz test, the Knüppel test, and the Kolmogorov-Smirnov test, which we use to test our hypothesis. Section IV describes our data set. In Section V, we present the empirical results of our study. Reasons for why the Ross recovery theorem empirically fails are given in Section VI. Section VII provides several robustness checks, while Section VIII concludes.

II. The Ross Recovery Theorem

An application of the recovery theorem requires the following assumptions: **(i)** The transition state prices $\pi_{i,j}$ follow a time-homogeneous process, which means that they are independent of calender time, **(ii)** the transition state prices need to be strictly positive, and **(iii)** the corresponding pricing kernel $m_{i,j}$ is transition independent, which means that it can be written as:

$$m_{i,j} = \delta \frac{u'_j}{u'_i} \tag{2}$$

for a positive constant δ and positive state-dependent values u'_j and u'_i. Ross (2015) suggests a possible interpretation of those values u' as marginal utilities, while viewing δ as a utility discount factor.

With this structure for the pricing kernel, the physical transition probabilities $p_{i,j}$ have the form:

$$p_{i,j} = \frac{\pi_{i,j}}{m_{i,j}} = \frac{1}{\delta} \cdot \frac{\pi_{i,j} \cdot u'_i}{u'_j}. \tag{3}$$

The Ross recovery theorem then allows to uniquely determine δ, all the u'_i, and the physical transition probabilities $p_{i,j}$ from the transition state prices $\pi_{i,j}$.

We illustrate the recovery theorem in a simple example with two states, state 0 and state 1. For any of the two possible initial states, the physical transition probabilities have to sum up to one:

$$\begin{aligned} p_{0,0} + p_{0,1} = 1 \quad &\Leftrightarrow \quad \frac{1}{\delta} \cdot \pi_{0,0} \cdot \frac{u'_0}{u'_0} + \frac{1}{\delta} \cdot \pi_{0,1} \cdot \frac{u'_0}{u'_1} = 1, \\ p_{1,0} + p_{1,1} = 1 \quad &\Leftrightarrow \quad \frac{1}{\delta} \cdot \pi_{1,0} \cdot \frac{u'_1}{u'_0} + \frac{1}{\delta} \cdot \pi_{1,1} \cdot \frac{u'_1}{u'_1} = 1. \end{aligned} \tag{4}$$

We can rewrite this system of equations in matrix form and obtain the following eigenvalue problem:

$$\begin{pmatrix} \pi_{0,0} & \pi_{0,1} \\ \pi_{1,0} & \pi_{1,1} \end{pmatrix} \cdot \begin{pmatrix} z_0 \\ z_1 \end{pmatrix} = \delta \cdot \begin{pmatrix} z_0 \\ z_1 \end{pmatrix} \quad \text{where} \quad z_0 = \frac{1}{u'_0} \quad \text{and} \quad z_1 = \frac{1}{u'_1}. \tag{5}$$

Given assumptions **(i)**, **(ii)**, and **(iii)**, an application of the Perron-Frobenius theorem leads to the result that there is only one eigenvector z with strictly positive entries z_0 and z_1. That eigenvector corresponds to the largest (and positive) eigenvalue δ of the eigenvalue problem. This property implies a unique positive pricing kernel $m_{i,j}$ as in Equation 2 and unique physical transition probabilities $p_{i,j}$ for $i, j = 0, 1$.

In the general setting with N different states, Ross defines the state transition matrix Π with entries $\pi_{i,j}$. Each row i in Π represents state prices of moving from a particular state i to any other state j. We always label the current state with $i = 0$ out of a set $I = \{-N_{low}, ..., 0, ..., N_{high}\}$ where $N = N_{low} + N_{high} + 1$. The ending transition state j is drawn from the same set I. The $0 - th$ row of Π contains the one period transition state prices, starting from the current state, which coincide with the one period spot state prices $\pi_{0,j}$. Analogous to the two state example, we solve the following N-dimensional eigenvalue problem:

$$\Pi z = \delta z, \quad \text{where} \quad z_i = \frac{1}{u_i'}. \tag{6}$$

Once we solve for z_i and δ, we are able to recover the physical transition probabilities $p_{i,j}$ from Equation 3.

III. Methodology

The basic ingredient missing at this point is the matrix Π of transition state prices, which are not readily observable in the market. Yet, transition state prices link spot state prices at different maturities with each other. The spot state prices $\pi^t_{0,i}$ (i.e., the subset of transition state prices that start at the current state 0 and have time to maturity t) can be readily obtained from observed option prices. We next detail our method for finding spot state prices before returning to the task of finding transition state prices from spot state prices. Another direct method even obviates the need for estimating the transition state price matrix altogether. We finally introduce our statistical tests.

A. Obtaining spot state prices from observed option prices

We collect European put- and call options quotes on the S&P 500. We average bid and ask quotes to obtain midpoint option prices, which we then transform to implied volatilities.

Quotes are only available for specific moneyness levels and maturities, yet we would like to obtain state prices on a grid compatible with the recovery theorem, which typically requires different levels of moneyness and different maturities. Thus, we generate a smooth implied volatility surface on a fine auxiliary grid, from which we later interpolate to the grid required for the recovery theorem.

We start with the fast and stable method of Jackwerth (2004), which finds smooth implied volatilities σ_i on a fine grid of states i for a fixed maturity. The fast and stable method minimizes the sum of squared second derivatives of implied volatilities (insuring smoothness of the volatility smile) plus the sum of squared differences between model and observed implied volatilities (insuring the fit to the options data) using the trade-off parameter λ.

We extend the method to volatility surfaces by adding a maturity dimension to the S&P 500 level dimension. Again, we minimize the sum of squared local total second implied volatility derivatives $\sigma''_{i,t}$ (insuring smoothness of the volatility surface and not only of the volatility smile) plus the sum of squared deviations of the model from the observed implied volatilities (insuring the fit of the surface) by using the trade-off parameter λ. We weight the squared second derivatives $(\sigma''_{i,t})^2$ with maturity t to compensate for the stronger curvature of the short-term volatility smile. The optimization problem is:

$$\begin{aligned} \min_{\sigma_{i,t}} \quad & \frac{1}{TN} \cdot \sum_{t=1}^{T} \sum_{i \in I} (\sigma''_{i,t})^2 \cdot t \quad + \quad \lambda \cdot \frac{1}{L} \cdot \sum_{l=1}^{L} \left(\sigma_{i(l),t(l)} - \sigma^{\text{obs}}_{i(l),t(l)} \right)^2 \\ & s.t. \\ & \sigma_{i,t} \ \geq \ 0, \end{aligned} \tag{7}$$

where $\sigma^{\text{obs}}_{i(l),t(l)}$ is the $l-th$ observed implied volatility (out of L observations) and I is a fine set of indexes for states i with a total number of N states. For the state space dimension, we discretize option strike prices with a step size of \$5. We then convert strike prices into moneyness levels by normalizing them with the current level of the S&P 500 index. We set sufficient upper and lower bounds for the moneyness such that all possible states with a non-zero probability of occurrence are covered.[11] We discretize the maturity with ten steps per month with a maximum maturity of twelve months, which gives us a total number of 120 maturity steps. This discretization insures that all available option prices lie on our fine

[11]Note that transition state prices in the discrete Ross recovery setting are only defined on a bounded state space. We argue that for our implementation these bounds are wide enough, as they, on each date, cover all states that are relevant to explain spot state prices with all maturities up to twelve months.

grid. We then solve Equation 7 to obtain the implied volatility surface on the fine grid. We provide more details in Appendix.A.

To obtain state prices on the coarser grid suitable for the recovery theorem, we linearly interpolate the fine implied volatility surface. From the implied volatilities on the coarser grid, we compute call option prices on the coarser grid and apply the Breeden and Litzenberger (1978) approach to find the spot state prices. Namely, for each maturity t, we take the numerical second derivative of the call prices to obtain spot state prices.[12]

B. Finding transition state prices from spot state prices

Now that we have the spot state prices in place, we return to the task of finding transition state prices from spot state prices. Following Ross (2015), we can identify the transition state prices $\pi_{i,j}$ since they link spot state prices at different maturities with one another. Using monthly transition state prices and spot state prices with monthly maturities of up to one year, we have the following relations:

$$\pi_{0,j}^{t+1} = \sum_{i \in I} \pi_{0,i}^{t} \cdot \pi_{i,j} \quad \forall j \in I, t = 0, ..., 11, \tag{8}$$

where today's spot state prices with a maturity of zero ($\pi_{0,i}^{0}$) are 0 for all states but the current state, for which the spot state price is one.

Equation 8 states that one can find the spot state price $\pi_{0,j}^{t+1}$ of reaching state j at maturity $t+1$ by adding up all the state price contributions of visiting state i one month earlier at maturity t ($\pi_{0,i}^{t}$) times the transition state price from i to j ($\pi_{i,j}$). We can exactly solve for all the transition state prices $\pi_{i,j}$, if the number of states N is equal to the number of (non-overlapping) transitions. With twelve transitions of one month each, we cover a whole year of maturities and can solve for only twelve states. Ross (2015) already mentions that such coarse non-overlapping grid leads to poorly discretized transition state prices and to coarse discrete physical probabilities based on the recovery theorem.[13]

[12] As we lose the lowest and the highest index level due to the numerical second derivative, we choose the initial coarse grid to be one index level too wide on either end. The spot state prices will then be exactly on the desired coarser grid.

[13] We use the non-overlapping approach on the twelve-by-twelve state space in the robustness section VII.A. As expected, the approach does not work and the realized future returns do not seem to be drawn from the recovered physical probability distribution.

Ross Basic

Instead of using the coarse non-overlapping grid of dimension twelve-by-twelve, we follow Audrino et al. (2015) and apply an *overlapping* approach to determining the transition state prices. Based on steps of one-tenth of a month (and a state price transition lasting one month, i.e., ten steps), our new relation is:

$$\pi_{0,j}^{t+10} = \sum_{i \in I} \pi_{0,i}^{t} \cdot \pi_{i,j} \quad \forall j \in I, t = 0, ..., 110. \tag{9}$$

This results in a total number of 111 overlapping transitions and thus allows N=111 states, which we choose to be equidistant and where we include the current state $i = 0$. Directly solving Equations 9 is not advisable as the problem is ill-conditioned. Rather, we impose an additional non-negativity constraint on the transition state prices $\pi_{i,j}$ and we back them out from the following least squares problem, which penalizes violations of Equations 9:

$$\min_{\pi_{i,j}} \quad \sum_{j \in I} \sum_{t=0}^{110} \left(\pi_{0,j}^{t+10} - \sum_{i \in I} \pi_{0,i}^{t} \cdot \pi_{i,j} \right)^2 \quad s.t. \quad \pi_{i,j} > 0. \tag{10}$$

We collect the transition state prices $\pi_{i,j}$ in the transition state price matrix Π and recover the matrix $P = [p_{i,j}]_{i,j \in I}$ of physical transition probabilities by applying the recovery theorem of Equation 6 to Π. We label this version of recovery **Ross Basic**.

We depict the results of our implementation for a typical day in our sample, February 17, 2010. Figure 1, Panel A, shows the interpolated smoothed implied volatility surface on the 111 by 111 state space. We note that the volatility surface is smooth yet passes through the observed implied volatilities (black squares). The volatility smile is clearly visible for short maturities and flattens out at longer maturities. Panel B shows the related spot state price surface, which also turns out to be smooth. Spot state prices tend to be high around the current state (moneyness of one) and are more spread out at larger maturities.

Figure 2, Panel A, illustrates the transition state prices, which best relate spot state prices at one maturity to those at a maturity one month later. We would expect large transition state prices on the main diagonal, as it is more likely to end up at states j which are close to the initial state i. However, the optimization quite often generates large state prices for states that are far away from the current state. This pattern of spurious transition state

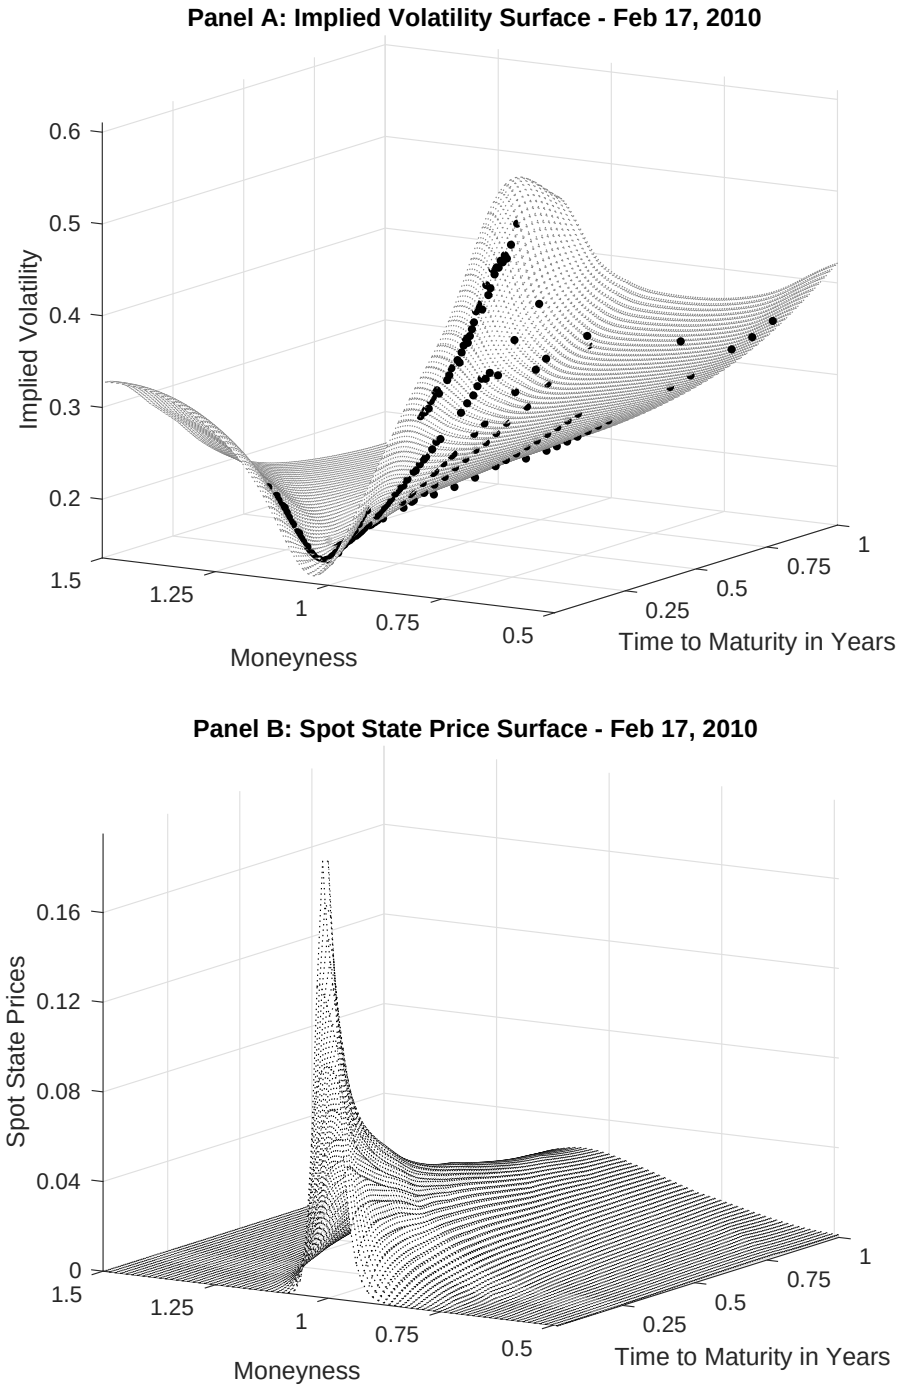

Figure 1. Implied volatility and spot state price surfaces. We show the interpolated implied volatility surface and the observed implied volatilities (as black dots) in Panel A. We show the related spot state price surface in Panel B. Based on data from February 17, 2010, we depict these surfaces for S&P 500 index options across maturities and moneyness levels.

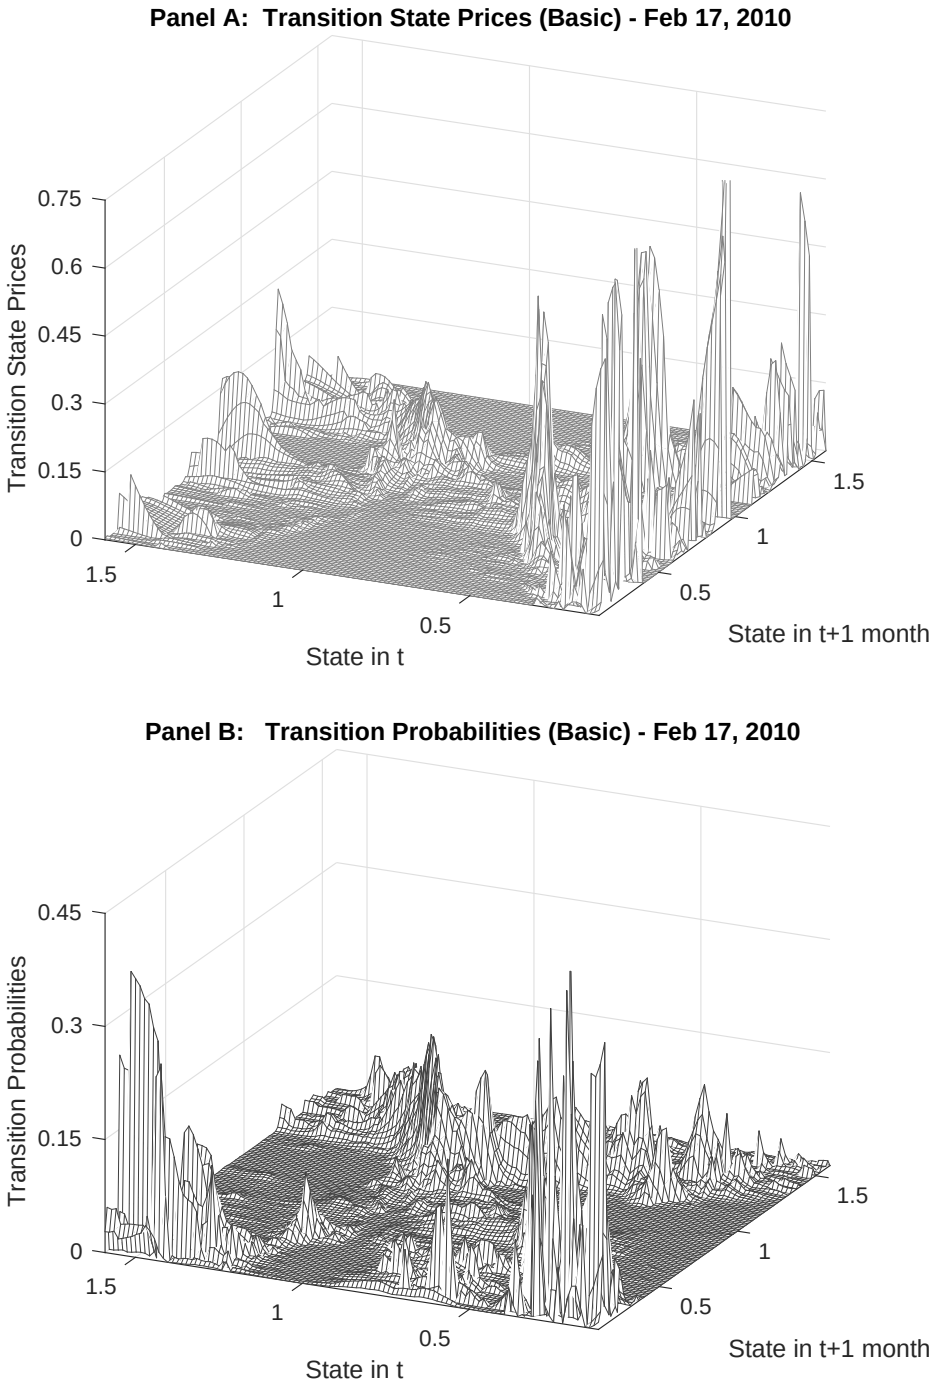

Figure 2. Transition state prices and recovered physical transition probabilities, Ross Basic. We show the transition state prices in Panel A and the corresponding recovered transition probabilities in Panel B as identified by the Ross Basic approach. All data are from February 17, 2010.

prices away from the main diagonal carries over to some extent to the recovered physical transition probabilities in Panel B of Figure 2.

The large transition state prices away from the main diagonal can occur because short maturity option prices are hardly affected by such irrelevant transition state prices, which link states that are not important for the short maturity spot state prices and thus the value of the short maturity options. The optimization can thus allocate mass to these irrelevant states to minimize the objective function, while not changing the short maturity spot state prices much in the process. As a result, some rowsums in Π have values much higher than one, which would imply high negative risk-free rates for some initial states. On our sample day, February 17, 2010, we observe that one-fourth of all one-month state-dependent risk-free rates are lower than -50% with the lowest value being -92% (-100% annualized), and that one-fourth of the rates are higher than 30% with the highest value being 165% (11,993,526% annualized). As the problem is very ill-conditioned, we impose further economic restrictions.

Ross Bounded

Namely, we first demand all rowsums of Π to lie in the interval $[0.9, 1]$. As the inverse of the rowsum is equal to one plus the risk-free rate for this state, we limit the monthly risk-free rates to between 0% and 11.11% (0% and 254.07% annualized). We again solve Equation 10 but additionally restrict the rowsums, before continuing with the recovery theorem. We label this version **Ross Bounded**.

Figure 3 illustrates the transition state prices in Panel A and the corresponding recovered physical transition probabilities in Panel B. Both surfaces are now highly concentrated around the current state. On the positive side, this eliminates high values in irrelevant states (i.e., far away from the main diagonal). Yet worryingly, even the values on the main diagonal fall off as we move away from the current state. The optimization allocates the restricted state prices in an almost uniform way for very low and very high states. As a result, we do not obtain the economically reasonable diagonal structure for the transition state prices with Ross Bounded.

Ross Unimodal

Next, we outright force the rows in our Π matrix to be unimodal with maximal values on the main diagonal. Once more, we solve Equation 10 but add the requirement of unimodality and that all rowsums of Π lie in the interval $[0.9, 1]$. We then proceed with the recovery theorem and label this version **Ross Unimodal**.[14]

[14]Related to our Ross Unimodal approach, Massacci et al. (2016) extract a 11 by 11 transition state price matrix from intraday S&P 500 option data and from Apple stock option data, respectively, and also force

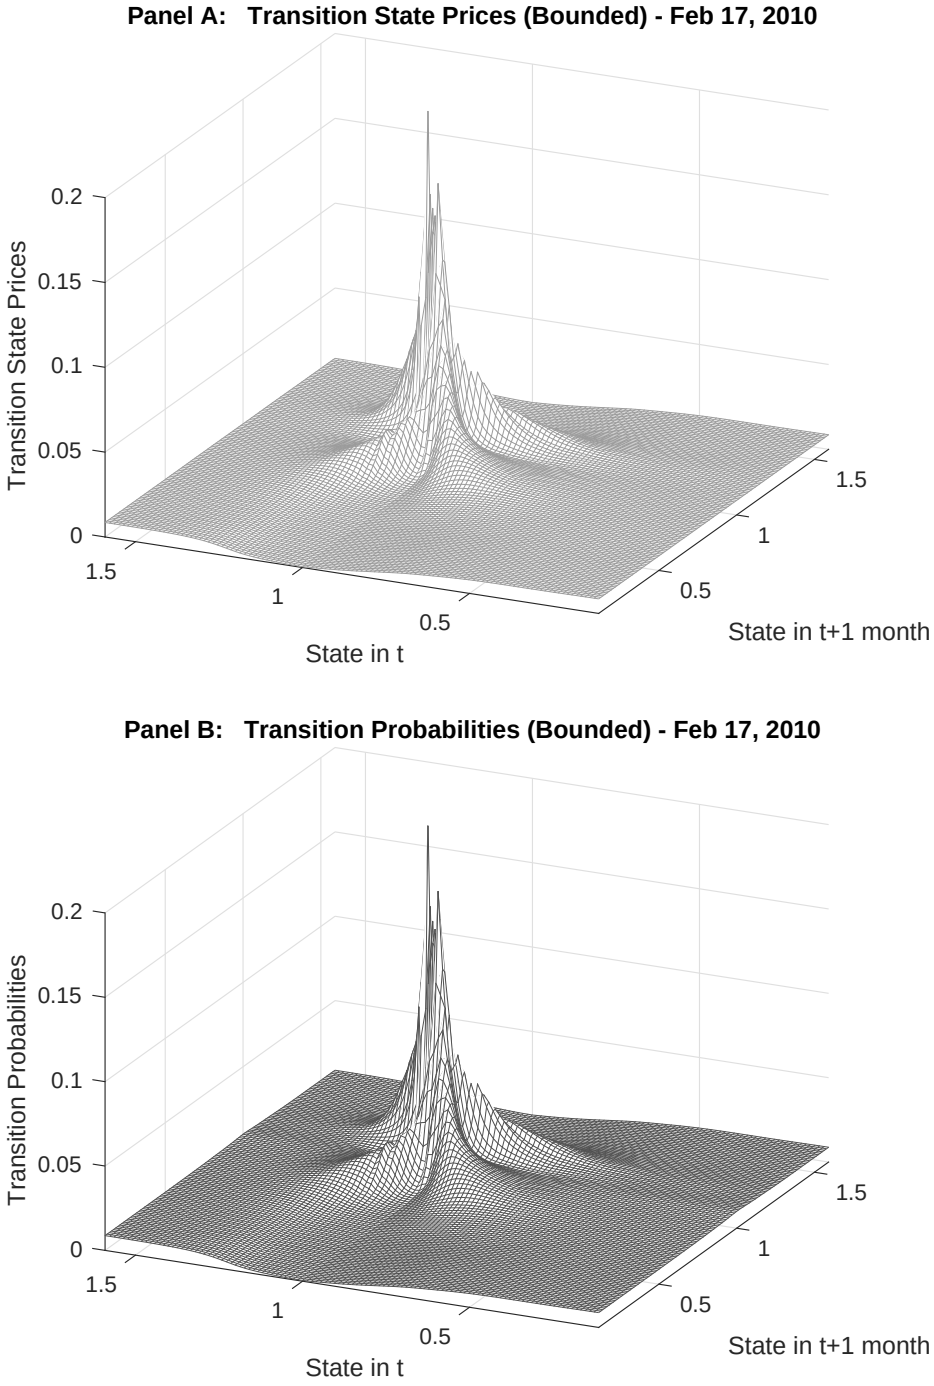

Figure 3. Transition state prices and recovered physical transition probabilities, Ross Bounded. We show the transition state prices in Panel A and the corresponding recovered transition probabilities in Panel B as identified by the Ross Bounded approach. All data are from February 17, 2010.

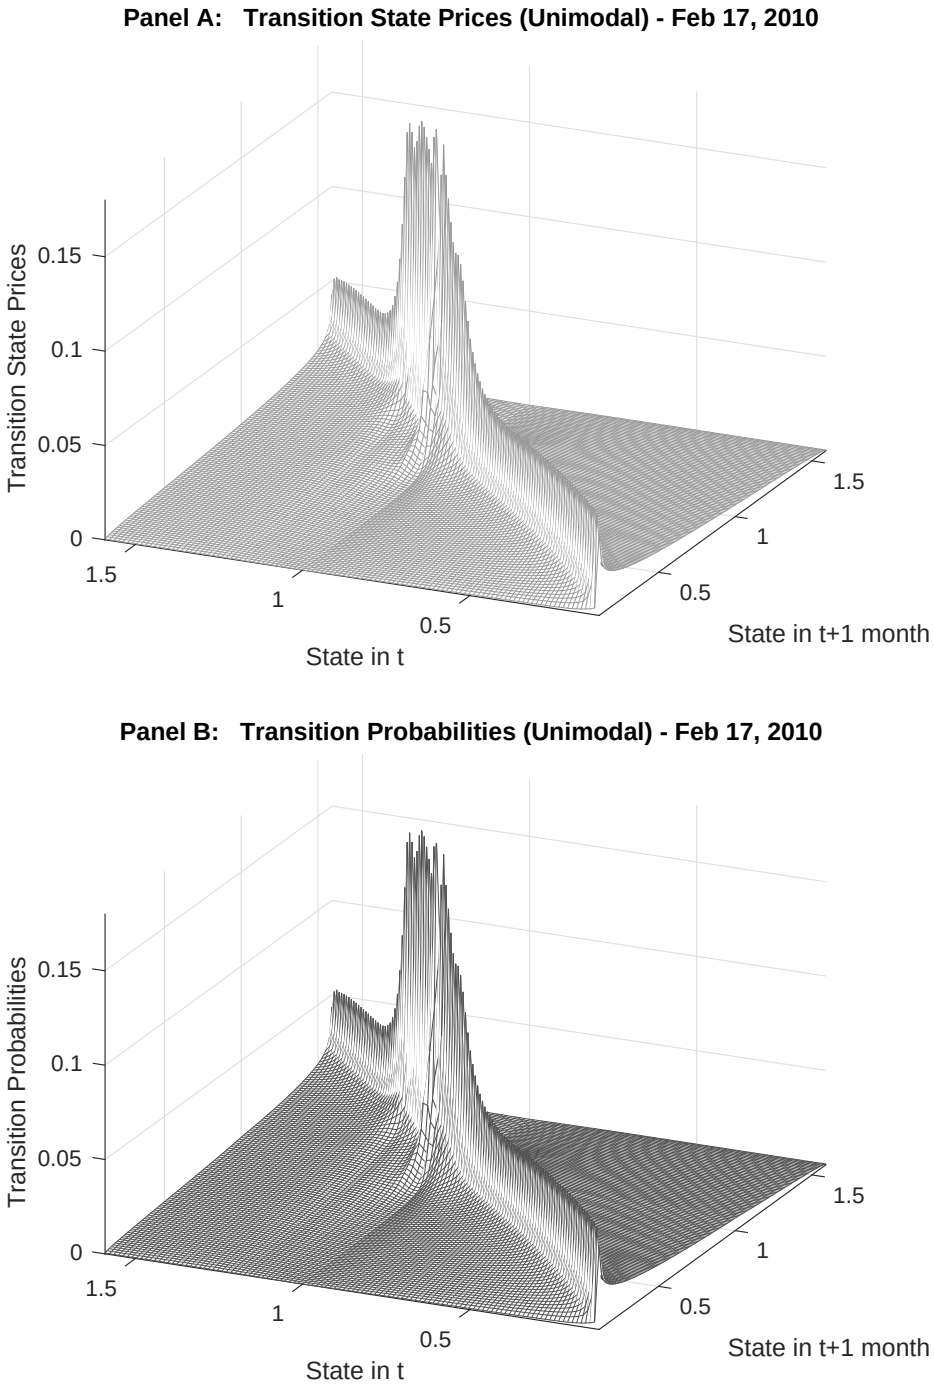

Figure 4. Transition state prices and recovered physical transition probabilities, Ross Unimodal. We show the transition state prices in Panel A and the corresponding recovered transition probabilities in Panel B as identified by the Ross Unimodal approach. All data are from February 17, 2010.

Figure 4 shows the transition state prices in Panel A and the corresponding recovered physical transition probabilities in Panel B for Ross Unimodal. By construction, the highest values line the main diagonal, steeply falling off further away from the main diagonal. However, higher values are again concentrated around the current state.

C. Recovery without using transition state prices, Ross Stable

The computation of transition state prices is a key challenge in the application of the recovery theorem. Yet, the one row in the transition state price matrix Π associated with the current state $i = 0$ offers a novel way out. For this current state, the transition state prices ought to coincide with the one period spot state prices, which we readily obtain from option prices. We use this insight to suggest an alternative recovery approach that does not require explicitly solving for the transition state prices.[15] The trick is using the eigenvalue problem in the recovery theorem as in Equation 6 and multiplying both sides from the left with the transition state price matrix Π:

$$\Pi \cdot \Pi z = \Pi \cdot \delta z = \delta(\Pi z) = \delta^2 z. \tag{11}$$

Again, the row of Π^2 associated with the current state $i = 0$ contains the spot state prices, but now with a maturity of two transition periods. In this case, the discount factor δ appears in the second power to account for the two periods. Iterating, we obtain the following relation:

$$\Pi^t z = \delta^t z \text{ with } t = 1, ..., T, \tag{12}$$

where t determines how often we apply the transition. For each t, we focus on the row in Π^t associated with the current state $i = 0$, where the t-period transition state prices coincide with the t-period spot state prices. We collect all those current rows with different maturities t. Full identification requires at least as many equations for different maturities t as there are number of states N, which results in the following system of equations (see Appendix.B for details):

transition state price matrices to have unimodal rows with maximum values on the main diagonal.

[15] See the independent derivation in Jensen et al. (2017).

$$\begin{pmatrix} \pi^1_{0,-N_{low}} & \pi^1_{0,-N_{low}+1} & \cdots & \pi^1_{0,N_{high}-1} & \pi^1_{0,N_{high}} \\ \pi^2_{0,-N_{low}} & \pi^2_{0,-N_{low}+1} & \cdots & \pi^2_{0,N_{high}-1} & \pi^2_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{T-1}_{0,-N_{low}} & \pi^{T-1}_{0,-N_{low}+1} & \cdots & \pi^{T-1}_{0,N_{high}-1} & \pi^{T-1}_{0,N_{high}} \\ \pi^{T}_{0,-N_{low}} & \pi^{T}_{0,-N_{low}+1} & \cdots & \pi^{T}_{0,N_{high}-1} & \pi^{T}_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} \frac{z_{-N_{low}}}{z_0} \\ \vdots \\ \frac{z_{-1}}{z_0} \\ 1 \\ \frac{z_1}{z_0} \\ \vdots \\ \frac{z_{N_{high}}}{z_0} \end{pmatrix} = \begin{pmatrix} \delta \\ \delta^2 \\ \vdots \\ \delta^{T-1} \\ \delta^T \end{pmatrix}. \tag{13}$$

We are worried that the system of equations is ill-conditioned and, thus, might violate sensible economic constraints. Namely, we want to insure that the utility discount factor δ and the resulting pricing kernel are non-negative. Thus, we penalize deviations from Equation 13 and include the two new constraints:

$$\min_{\left[\frac{z_j}{z_0}\right],\delta} \quad \sum_{t=1}^{T} \left(\sum_{j \in I} \pi^t_{0,j} \cdot \left[\frac{z_j}{z_0}\right] - \delta^t \right)^2 \quad s.t. \quad \left[\frac{z_j}{z_0}\right] > 0,\ 1 > \delta > 0. \tag{14}$$

We still need to decide on the length of the transition period. Using twelve non-overlapping periods of one month each gives us only twelve transitions and allows for at most twelve states. This results in too coarse a grid.[16] Instead, we again use 120 periods of one-tenth of a month each. Here, we make use of the property that the structure of the pricing kernel in the setting of Ross (2015) remains the same for different maturities and only varies by a factor:[17]

$$m^t_{0,j} = \delta^{t-1} \cdot m_{0,j}, \tag{15}$$

where $m^t_{0,j}$ is the *spot* pricing kernel with a maturity of t transition periods and $m_{0,j}$ is the *spot* pricing kernel with a maturity of one transition period. This allows us to use a fine moneyness grid defined on 120 points, corresponding to 120 maturities that are spaced

[16] We still use the coarse grid in our robustness checks, Section VII.A, but, as expected, future realized returns are incompatible with the recovered physical probability distribution.

[17] See Appendix.B for details.

one-tenth of a month apart. Once we solve Equation 14 for the one period pricing kernel, we use Equation 15 to find the ten period (i.e., one-month) pricing kernel. We use this one-month pricing kernel to transform one-month spot state prices into one-month physical probabilities.[18] We label this approach **Ross Stable**. Note that we cannot provide the corresponding figures for the transition state prices and the transition physical probabilities as we no longer compute them explicitly.

D. Testing the recovered physical probabilities

We finally have the physical probabilities, based on our four version of Ross recovery, in hand and want to investigate how good these forecasts are, which are solely based on the option prices. We will shortly add two further competing approaches, namely Power Utility (the physical distribution implied by a power utility pricing kernel in combination with the spot state prices) and Historical Return Distribution (the physical distribution based on the past five years of monthly returns). Our hypothesis is that:

H0: Future realized monthly S&P 500 returns are drawn from the recovered physical distribution.

We test our hypothesis as follows. Each month, we find the date τ which is 30 calendar days before the option expiration date. We record the realized future return (i.e., the return from date τ to the expiration date) on the S&P 500 and label it R_τ. That return is one realization drawn from the true physical distribution p_τ. Next, we turn to the recovery theorem. We recover the physical spot distribution $\hat{p}_\tau$ and the corresponding cumulative distribution $\hat{P}_\tau$ for date τ. We can find the percentile of the recovered cumulative distribution $\hat{P}_\tau$ that corresponds to the realized return R_τ and we collect those percentiles x_τ for all dates.[19] Under the assumption that the recovered distribution is indeed the one from which the return was drawn (i.e., $\hat{p}_\tau = p_\tau$), the percentiles should be i.i.d. uniformly distributed.[20]

[18] As we solve for the pricing kernel with the least squares approach of Equation 14, the system of Equations 13 does not hold exactly. As a result, the recovered physical spot probabilities do not necessarily sum up to one, and so we normalize them.

[19] The values of the recovered cumulative distributions $\hat{P}_\tau$ lie on a discrete grid, whereas the future realized returns typically do not lie on the corresponding grid points. We therefore interpolate the recovered cumulative distribution linearly to obtain the transformed points x_τ.

[20] See Bliss and Panigirtzoglou (2004) and Cuesdeanu and Jackwerth (2016).

To check on the uniformity of the percentiles x_τ, we use three different tests: The Berkowitz test, the Knüppel test, and the Kolmogorov-Smirnov test, see Appendix.C for details. The Berkowitz test has been used in Bliss and Panigirtzoglou (2004). In that study, the authors argue that it is superior to the Kolmogorov-Smirnov test in small samples with autocorrelated data. While being more powerful than the Kolmogorov-Smirnov test, the Berkowitz test bases its conclusion about standard normality on just the first and the second moments, while ignoring higher moments. The Knüppel (2015) test has the advantage of testing for higher moments, can deal with autocorrelated data, and still has enough power in small samples.[21]

We apply our three tests of uniformity in combination with all recovery versions (Ross Basic, Ross Bounded, Ross Unimodal, and Ross Stable) to see if the recovered physical distributions are compatible with future realized returns. For comparison, we use two additional benchmark models for the physical distribution. For one, we use the empirical cumulative distribution of the past five years of monthly S&P 500 returns, labeled **Historical Return Distribution**. For the other, we assume a representative investor having a power utility with a risk aversion coefficient of four.[22] Based on the associated pricing kernel, we transform the spot state prices into a physical distribution. For comparability, we use the same one-month spot state prices as in Ross Stable, which lie on a moneyness grid defined on 120 points. We label this approach **Power Utility** with $\gamma = 4$.

IV. Data

We use OptionMetrics to obtain end-of-day option data on the S&P 500 index. We use the midpoint of bid and ask option quotes as our option prices. Consistent with the literature, we only use out-of-the-money put- and call options with a positive trading volume and eliminate all options that violate no arbitrage constraints. We fit the implied dividend yield according to put-call parity where the risk free rate is given by the interpolated zero-curve from OptionMetrics. We consider monthly dates τ, which we find by going 30 calendar days back in time from the expiration date. We end up with 223 recovered distributions. Our sample period, as well as our option data, ranges from January 1996 to August 2014.

[21] We do not consider the Cramer-van-Mises test as it is an integrated version of the Kolmogorov-Smirnov test and yields very similar results.

[22] The results in Bliss and Panigirtzoglou (2004) suggest that a risk aversion factor of four in a power utility framework is a reasonable choice.

For the historical return distribution, we further obtain end-of-day S&P 500 index levels from Datastream. We compute S&P 500 returns from January 1991 to August 2014, which includes a five year period prior to January 1996. Finally, we collect prices of the 30-year bond futures from Datastream, which we use for an additional analysis in Section VI.C.

V. Empirical Results

We are now ready to investigate our central question. Are future realized S&P 500 returns drawn from recovered physical probabilities? The sobering answer is in Table I. All four versions of the recovery theorem (Ross Basic, Ross Bounded, Ross Unimodal, and Ross Stable) strongly reject our null hypothesis (p-values less that 0.028) for all three tests, which we employ (Berkowitz, Knüppel, and Kolmogorov-Smirnov).

In contrast, our simple benchmark models (Power Utility with $\gamma = 4$ and Historical Return Distribution) are not rejected by any of our three tests (p-values of more than 0.294). We are thus facing a complete empirical failure of the recovery theorem, while a power utility setting or even a five year distribution of historical returns cannot be rejected by the data.

To better understand what drives our results, we take a closer look at the recovered probabilities in Figure 5. We start our discussion with Power Utility in Panel E, as it "works" in explaining the data and is particularly simple. The black line shows the spot state prices, which are derived from the option prices and are the same in all six panels.[23] Note that plotting the risk-neutral distributions would look just the same as plotting the spot state prices, as the one-month risk-free rate adjustment is invisible in the figures. The Power Utility approach changes the spot state prices into physical probabilities (dashed gray), for which we cannot reject our hypothesis that future realized returns are drawn from it. A successful method thus needs to right-shift the physical probabilities to be compatible with the data (implying a positive market-risk premium).

The Historical Return Distribution in Panel F also works in that our main hypothesis cannot be rejected. The physical distribution is the kernel-smoothed histogram of the one-month non-overlapping S&P 500 returns from the past five years.[24] The physical distribution

[23]There is a tiny difference in how they plot as we have 111 states in Panels A-C and we have 120 states in Panels D and E.

[24]The kernel smoothing uses Matlab's ksdensity with the default bandwidth, see Bowman and Azzalini (1997). New York: Oxford University Press Inc., 1997. For plotting, we interpolate the kernel smoothed density onto 120 states that we use for Power Utility and Ross Stable and normalize the result such that it

Table I. **Tests of the recovered physical probabilities.** We present our results if future realized returns are drawn from physical probabilities generated by one of our six approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, Power Utility, and Historical Return Distribution. For each approach, we show the p-values from the Berkowitz, Knüppel, and Kolmogorov-Smirnov tests for uniformity of the percentiles of future realized returns under the model physical cumulative distribution.

Recovery Approach	**Berkowitz** p-value	**Knüppel** p-value	**Kolmogorov-Smirnov** p-value
Ross Basic $\pi_{i,j} > 0$	0.018	0.027	0.000
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.005	0.002	0.008
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.001	0.000	0.028
Ross Stable Do not use transition state prices	0.010	0.015	0.004
Power Utility with $\gamma = 4$	0.697	0.320	0.547
Historical Return Distribution	0.294	0.480	0.347

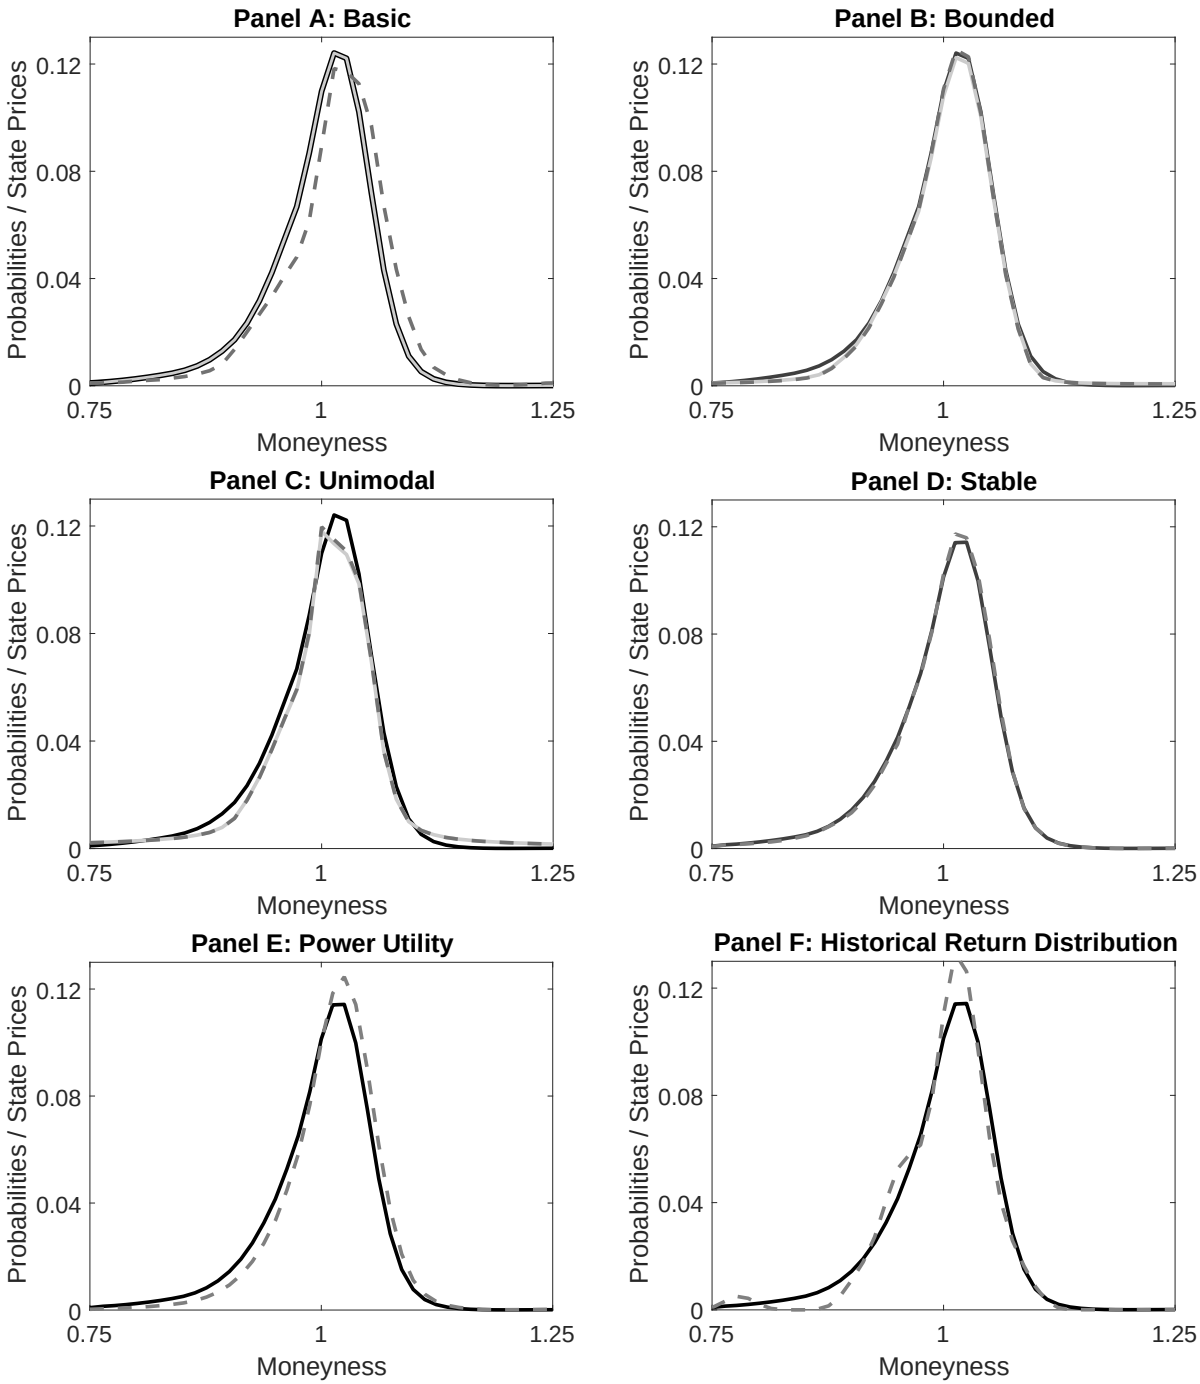

Figure 5. State prices and recovered physical probabilities. We depict spot state prices (black lines), transition state prices (light gray lines), and recovered physical probabilities (gray dashed lines) on February 17, 2010. Our methods are Ross Basic in Panel A, Ross Bounded in Panel B, and Ross Unimodal in Panel C, Ross Stable in Panel D, Power Utility with $\gamma = 4$ in Panel E, and the kernel-smoothed Historical Return Distribution in Panel F.

is less smooth than in Panel E, as some of the jaggedness of the histogram is still present in the kernel distribution.

We now turn to the recovery approaches (Panels A-C) that have an additional light gray line for the transition state prices. This is because the models allow for the transition state prices to differ from the spot state prices, thus incorporating a pricing error for the option prices. Yet, Ross Basic in Panel A does not use this possibility, which is why the spot and transition state prices plot on top of each other. The reason can be found in the lack of economic constraints on the transition state prices (other than positivity), which then allows the optimization to closely follow the spot state prices. As a result, Ross Basic exhibits implausible fluctuations for the transition state prices and rowsums, which imply extremely large negative or positive monthly risk-free rates, ranging from -92% to 165% on a typical sample day (February 17, 2010). The physical distribution is somewhat right-shifted but insufficiently so, as we reject our main hypothesis for Ross Basic as well as for all other recovery approaches.

Adding economic constraints on the rowsums in Ross Bounded (Panel B) and unimodality in Ross Unimodal (Panel C) leads to a slight separation of the transition state prices from the spot state prices. We quantify this mispricing below in Section VI.D. Also, the physical probabilities are almost identical to the transition state prices. Chaining the results together, the physical distribution remains close to the transition state prices, which are close to the spot state prices. But then the recovered physical distribution ends up being too close to the spot state prices, and we reject our main hypothesis that the recovered distribution is compatible with the future realized returns. Finally, in Ross Stable, we do not explicitly compute transition state prices. The physical probabilities are again very close to the spot state prices, and we reject our main hypothesis again for Ross Stable.

Summing up, all recovery approaches, as opposed to our simple benchmark models, are incompatible with future realized S&P 500 returns. Ross Basic suggests extreme fluctuations in the transition state prices and the risk-free rates in different states. The other recovery approaches cannot generate a sufficiently high risk premium as the recovered physical distribution stays too close to the spot state prices.

---

sums to one.

VI. Reasons for Failure

We now investigate in more detail why the recovery theorem empirically fails. First, we analyze the recovered pricing kernels, uncover how those pricing kernels depend on transition state prices, and show that the recovered pricing kernels are often too flat. Second, switching to a time series perspective, we compare realizations of the recovered pricing kernel with theoretically motivated pricing kernels, namely, the minimum variance pricing kernel and the transitory pricing kernel component. Third, we look into the empirical implications of time-homogeneous transition state prices, which leads to poorly fitted longer-dated options. Fourth, we simulate an economy in which Ross recovery holds, complete with option prices and future realized returns. We then perturb the option prices and test if future returns are compatible with the recovered physical distribution based on the perturbed option prices.

A. Recovered Pricing Kernels

To better understand our empirical findings, we first analyze our recovered pricing kernels. Yet how should the pricing kernel look across values of S&P 500 returns? From basic theory, we expect the pricing kernel to be positive and monotonically decreasing, and to reflect the behavior of a risk averse representative investor. Jackwerth (2004), Ait-Sahalia and Lo (2000), and Rosenberg and Engle (2002), however, find that the empirical pricing kernel is locally increasing, a behavior that is referred to as the pricing kernel puzzle. In these papers, physical probabilities are backed out from past S&P 500 index returns, while Ross recovery implies pricing kernels based on forward-looking information. However, Cuesdeanu and Jackwerth (2016) confirm the existence of the prizing kernel puzzle in forward-looking data. Thus, we might expect the recovered pricing kernels to be monotonically decreasing (standard case) or to be locally increasing (pricing kernel puzzle case).

Figure 6 shows the pricing kernels obtained from different recovery approaches on February 17, 2010. The black line shows the *implied* pricing kernel for each approach, measured as *spot* state prices divided by the recovered physical probabilities. We start with Power Utility in Panel E, as the pricing kernel is, by construction, monotonically decreasing and well-grounded theoretically. From our main result, we also know that this pricing kernel translates the spot state prices into physical probabilities that are compatible with future realized returns. In a way, the power pricing kernel "works", while the pricing kernels based on the recovery theorem do not. We now try to understand why the latter do not work.

Most similar to the power pricing kernel is the pricing kernel for Ross Basic in Panel A. It

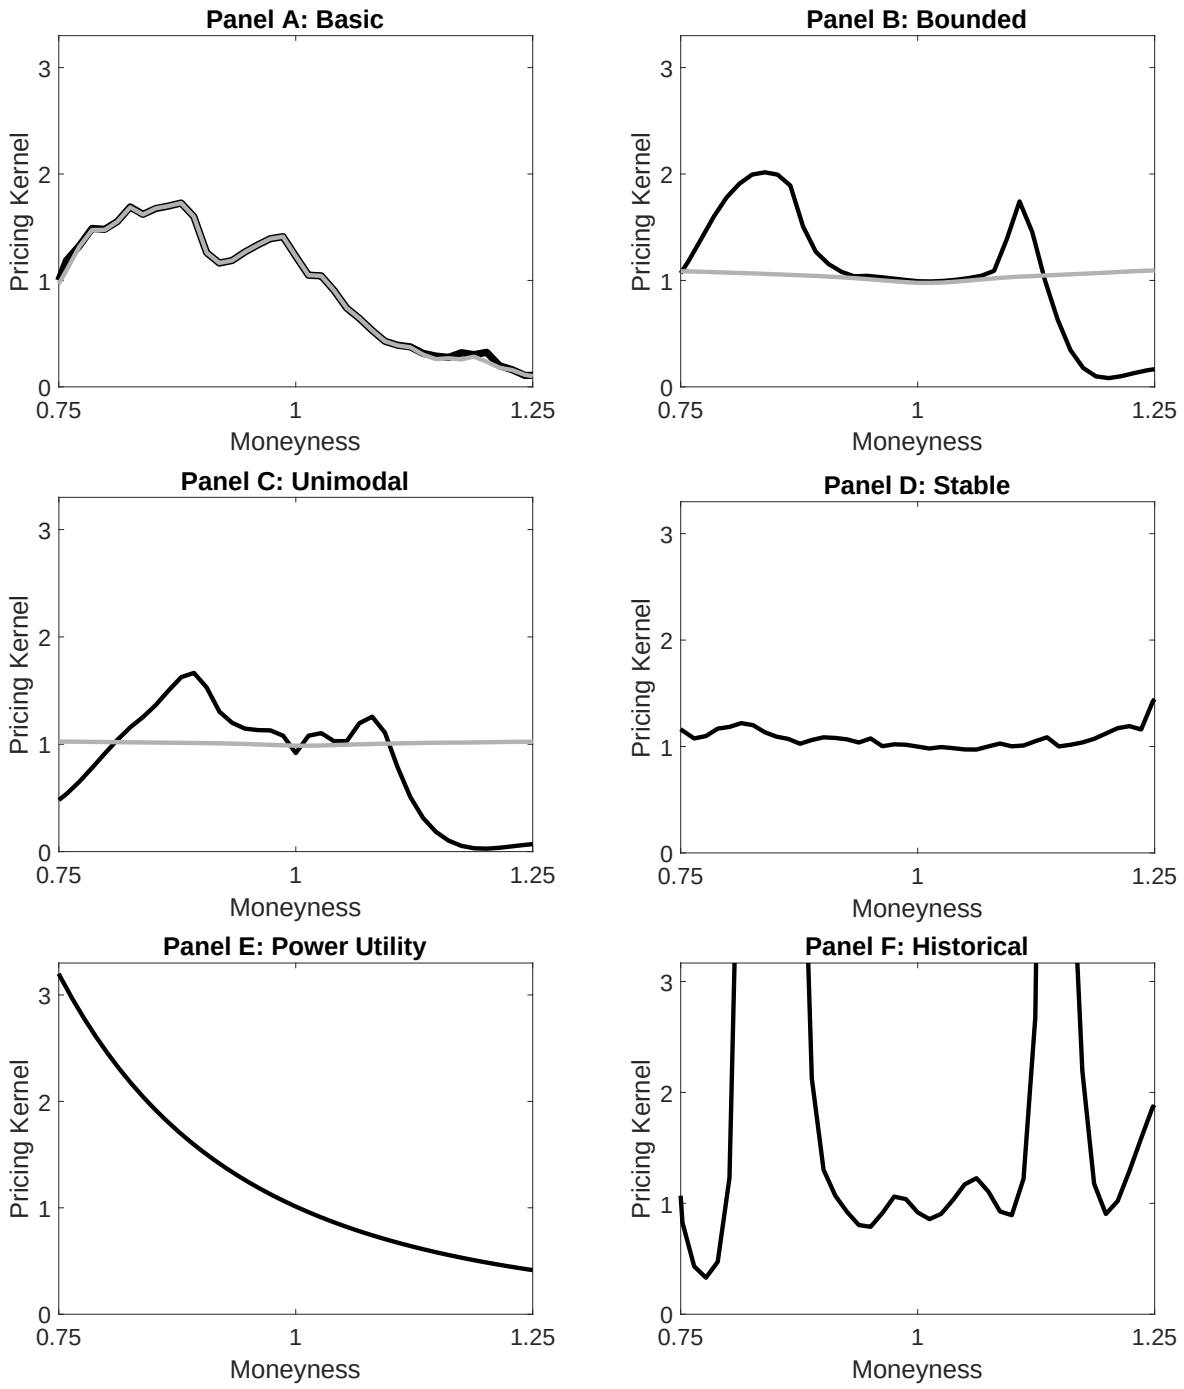

Figure 6. Pricing kernels. We present pricing kernels for different recovery approaches on February 17, 2010. Black lines depict *implied* pricing kernels, measured as spot state prices divided by recovered probabilities, while gray lines depict *model* pricing kernels measured as transition state prices for the current state divided by recovered probabilities. Panel A shows the pricing kernels for Ross Basic, Panel B for Ross Bounded, Panel C for Ross Unimodal, Panel D for Ross Stable (only the implied kernel exists), Panel E for Power Utility with $\gamma = 4$ (only the implied kernel exists), and Panel F for Historical Return Distribution (only the implied kernel exists).

is not very smooth but somewhat decreasing, yet less so than the power pricing kernel. As a result, the shift from state prices to physical probabilities is insufficient in that we reject our hypothesis that future realized returns are drawn from the recovered physical distribution. Here, we also depict as a gray line the *model* pricing kernel, measured as *transition* state prices for the current state divided by the recovered physical probabilities. Any difference in the two pricing kernels would be due to the optimization not being able to exactly match the spot state prices (and thus the observed option prices). For Ross Basic, this is not an issue as the optimization is free to fit option prices as long as the transition state prices are positive. We learned that this freedom comes at the cost of extreme rowsums, which in turn lead to extreme monthly risk-free rates, ranging from -92% to 165%.

Once we implement reasonable economic constraints in Ross Bounded and Ross Unimodal (Panels B and C), the implied pricing kernels become even more wavy overall and flatter for center moneyness levels of about 0.9 to 1.1. The recovered physical probabilities remain closer to the spot state prices and move further away from the distribution of future realized returns.[25] Interestingly now, the implied and the model pricing kernels fall apart, indicating that the optimization struggles to match the spot state prices as the transition state prices now need to satisfy our economic constraints. The model pricing kernel (the part driven by the recovery theorem and not due to the fit of option prices) is now virtually flat. The recovery theorem thus cannot generate a decreasing pricing kernel any more, when it is even gently constrained. Note that the requirement of monthly risk-free rates to lie between 0% and 11.11% is not very onerous.[26]

We shed further light on the flat pricing kernels once we realize that the rowsums of the transition state price matrix Π are strongly negatively related to the model pricing kernel.[27] Thus, for a decreasing pricing kernel (such as the power pricing kernel, which works well empirically), we would need increasing rowsums (i.e., decreasing interest rates) across states. Yet even our modest economic constraints on the rowsums render the model pricing kernels virtually flat. Ross recovery thus does not seem to be capable of generating non-flat model pricing kernels without unreasonably extreme risk-free rates.

[25]See e.g. Bliss and Panigirtzoglou (2004), who confirm established results in the literature, that risk-neutral probabilities perform worse in forecasting the distribution of future returns.

[26]The requirement of unimodality is also not very strong and only applies to Ross Unimodal in Panel C, but not to Ross Bounded in Panel B.

[27]Audrino et al. (2015) also noticed this relation.

We also find a flat implied pricing kernel for Ross Stable in Panel D.[28] Yet that flatness stems partially from the numerical implementation. We first recover a pricing kernel with maturity of 0.1 months to allow a larger number of 120 states (instead of only twelve), and then convert the 0.1-month pricing kernel into a one-month pricing kernel by the multiplicative adjustment of Equation 15. The negligible curvature of the 0.1-month pricing kernel then directly translates into an almost flat one-month pricing kernel, which again implies that we recover physical probabilities that are close to the spot state prices.

Last, the implied pricing kernel for Historical Return Distribution in Panel F is rather irregular on this particular day, even after we smoothed the historical distribution through a kernel density. Yet in general and across our whole sample, we cannot reject our main hypothesis and the implied pricing kernels manages to translate the spot state prices into generally right-shifted physical distributions.

We conclude that the major problem with Ross Basic are the extreme fluctuations of transition state prices and risk-free rates associated with the different states. As a result, the resulting pricing kernel is quite wavy as a function of moneyness. The major problem with the other Ross recovery approaches is their inability to generate sloped pricing kernels. All recovery approaches produce pricing kernels, which are incompatible with future realized returns.

B. Time Series of Pricing Kernels

As we have seen above for a particular day (February 17, 2010), the recovered pricing kernels are often rather flat. Yet, how should they look like in general? A time series perspective can help here and we ask for a time series of positive pricing kernel realizations m_τ being able to price the market and the risk-free asset.[29] Out of the many pricing kernels satisfying these constraints, we start with the Minimum Variance pricing kernel (below, we also look at the Maximum Entropy pricing kernel, see Ghosh et al. (2017)):

$$\min_{m_\tau} Var(m_\tau) \quad \text{s.t.} \quad \frac{1}{\mathcal{T}}\sum_{\tau=1}^{\mathcal{T}} m_\tau(R_\tau - Rf_\tau) = 0, \quad \frac{1}{\mathcal{T}}\sum_{\tau=1}^{\mathcal{T}} m_\tau Rf_\tau = 1, \quad m_\tau > 0. \tag{16}$$

[28]The model pricing kernel does not exist as we never explicitly compute the transition state prices.

[29]See, e.g., Cochrane (2000) page 6 for more details.

with R_τ being the realized 30-day future return market return and Rf_τ the realized 30-day future risk-free return at date τ.

We are curious how that Minimum Variance pricing kernel compares to a time series of realizations of the pricing kernel implied by a particular recovery approach. We determine such a pricing kernel realization for each date τ as the value of the method's implied pricing kernel that corresponds to the market return R_τ.[30]

To compare both time series, we regress the realized pricing kernel on the Minimum Variance pricing kernel:

$$m_\tau^{\text{Recovery}} = \beta_0 + \beta_1 m_\tau^{\text{MinVar}} + \epsilon_\tau, \quad \tau = 1, ..., \mathcal{T}. \tag{17}$$

If the recovered realized pricing kernel equals the Minimum Variance pricing kernel, we should obtain $\beta_0 = 0$ and $\beta_1 = 1$. However, we highlight that equality to the Minimum Variance pricing kernel is not required for a good pricing kernel candidate, as there are many other pricing kernels that are able to price the market and the risk-free asset. Since the Minimum Variance pricing kernel exhibits the lowest variation, we, in fact, would expect the slope β_1 to be larger than one and and the intercept β_0 to be smaller than zero in order to match the mean. Thus, we are curious about the coefficient of determination R^2 of Regression 17, which we view as a variance-invariant measure for similarity between the two time series.

Table II presents the regression results for all Ross recovery approaches, as well as for Power Utility. We do not include Historical Return Distribution here, as we already use the Minimum Variance pricing kernel as a return-based benchmark.

Among all approaches, Ross Basic implies the most extreme results with a highly negative but insignificant intercept of -12.71 and a highly positive but insignificant slope of 21.03. With an adjusted R^2 that is virtually zero, it also shows the least similarity to the Minimum Variance pricing kernel.

For Ross Bounded, Ross Unimodal, and Ross Stable, we already find more similarity, with an adjusted R^2 that ranges from 0.045 to 0.211. All three approaches imply positive intercepts and slopes that are smaller than one. One explanation is that their pricing kernels are close to risk-neutrality, which makes them less variable than the Minimum Variance

[30]We note that the resulting pricing kernel realizations are not constrained to exactly price the market and the risk-free asset. This is in contrast to the Minimum Variance pricing kernel in Equation 16.

Table II. **Realized Pricing Kernels.** We present the results from regressing the time series of realized pricing kernels for different recovery approaches and for Power Utility on the time series of the realized Minimum Variance pricing kernel. We do not present results for Historical Return Distribution as we already use the Minimum Variance pricing kernel as a return-based benchmark.

Recovery Approach	Intercept β_0	Slope β_1	adjusted R^2
Ross Basic $\pi_{i,j} > 0$	-12.71	21.03	-0.003
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.14	0.89***	0.211
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.38***	0.69***	0.117
Ross Stable Do not use transition state prices	0.80***	0.20***	0.045
Power Utility with $\gamma = 4$	-2.52***	3.56***	0.911

* indicates significance at 10%, ** indicates significance at 5%, *** indicates significance at 1%.

pricing kernel.[31] For Power Utility, we finally find a negative intercept and a slope that is larger than one. Further, the adjusted R^2 of 0.911 indicates a high similarity to the Minimum Variance pricing kernel.

How do those realized pricing kernels look like? Figure 7 depicts realized pricing kernels (gray squares) for our four Ross recovery approaches (Panel A-D), for Power Utility (Panel E) and for the Minimum Variance pricing kernel (Panel F). The black line in each panel represents a Gaussian kernel regression line through the realized pricing kernel values.

The realized pricing kernel for Ross Basic moves around widely and demonstrates the instability of that approach over time. In comparison, Ross Bounded, Ross Unimodal, and Ross Stable are less variable and exhibit rather flat pricing kernels. These cluster around one, approximating risk-neutrality. The picture is different for Power Utility and the Minimum Variance pricing kernel, which are both decreasing in future realized returns.

While Figure 6 in the previous section recovered pricing kernels for only one particular trading day (February 17, 2010), Figure 7 depicts the whole time series of realized pricing kernels. For Ross Bounded, Ross Unimodal, and Ross Stable, the shape of the realized pricing kernel time series is very similar to the shape of the implied pricing kernels on this particular trading day.

As an alternative to the Minimum Variance pricing kernel, we look at another valid pricing kernel based on Maximum Entropy. Here, we determine the realized pricing kernels, given the same constraints as in Equation 16, by minimizing the distance of the risk-neutral measure Q and the physical measure P with the Kullback-Leibler divergence $D_{KL}(Q||P)$.[32] The Maximum Entropy pricing kernel, however, turns out to be almost identical to the Minimum Variance pricing kernel and only exhibits a slightly higher variance. Going from Minimum Variance to Maximum Entropy does not significantly change our results.

We conclude that Ross Basic is so noisy that it hardly bears any resemblance with the Minimum Variance pricing kernel. On the other hand, even mild economic constraints in Ross Bounded, Ross Unimodal, or Ross Stable lead to pricing kernels, which are almost risk-neutral and, again, do not resemble the Minimum Variance pricing kernel. Only Power

[31] We note that, as these realized pricing kernels are not restricted to price the market and the risk-free asset, they are allowed to have a lower variation than the Minimum Variance pricing kernel.

[32] The Kullback-Leibler divergence is given by $\sum_{\tau=1}^{\mathcal{T}} q_\tau \dot{\log}\left(\frac{q_\tau}{p_\tau}\right)$ with risk neutral probabilities q_τ and physical probabilities p_τ. Using $\frac{q_\tau}{p_\tau} = m_\tau Rf_\tau$ and the assumption that all returns in the time series appear with equal probability ($p_\tau = \frac{1}{\mathcal{T}}$), we end up minimizing $\mathcal{T} \sum_{\tau=1}^{\mathcal{T}} m_\tau Rf_\tau \log(m_\tau Rf_\tau)$, where Rf_τ is the future 30-day risk-free return at date τ.

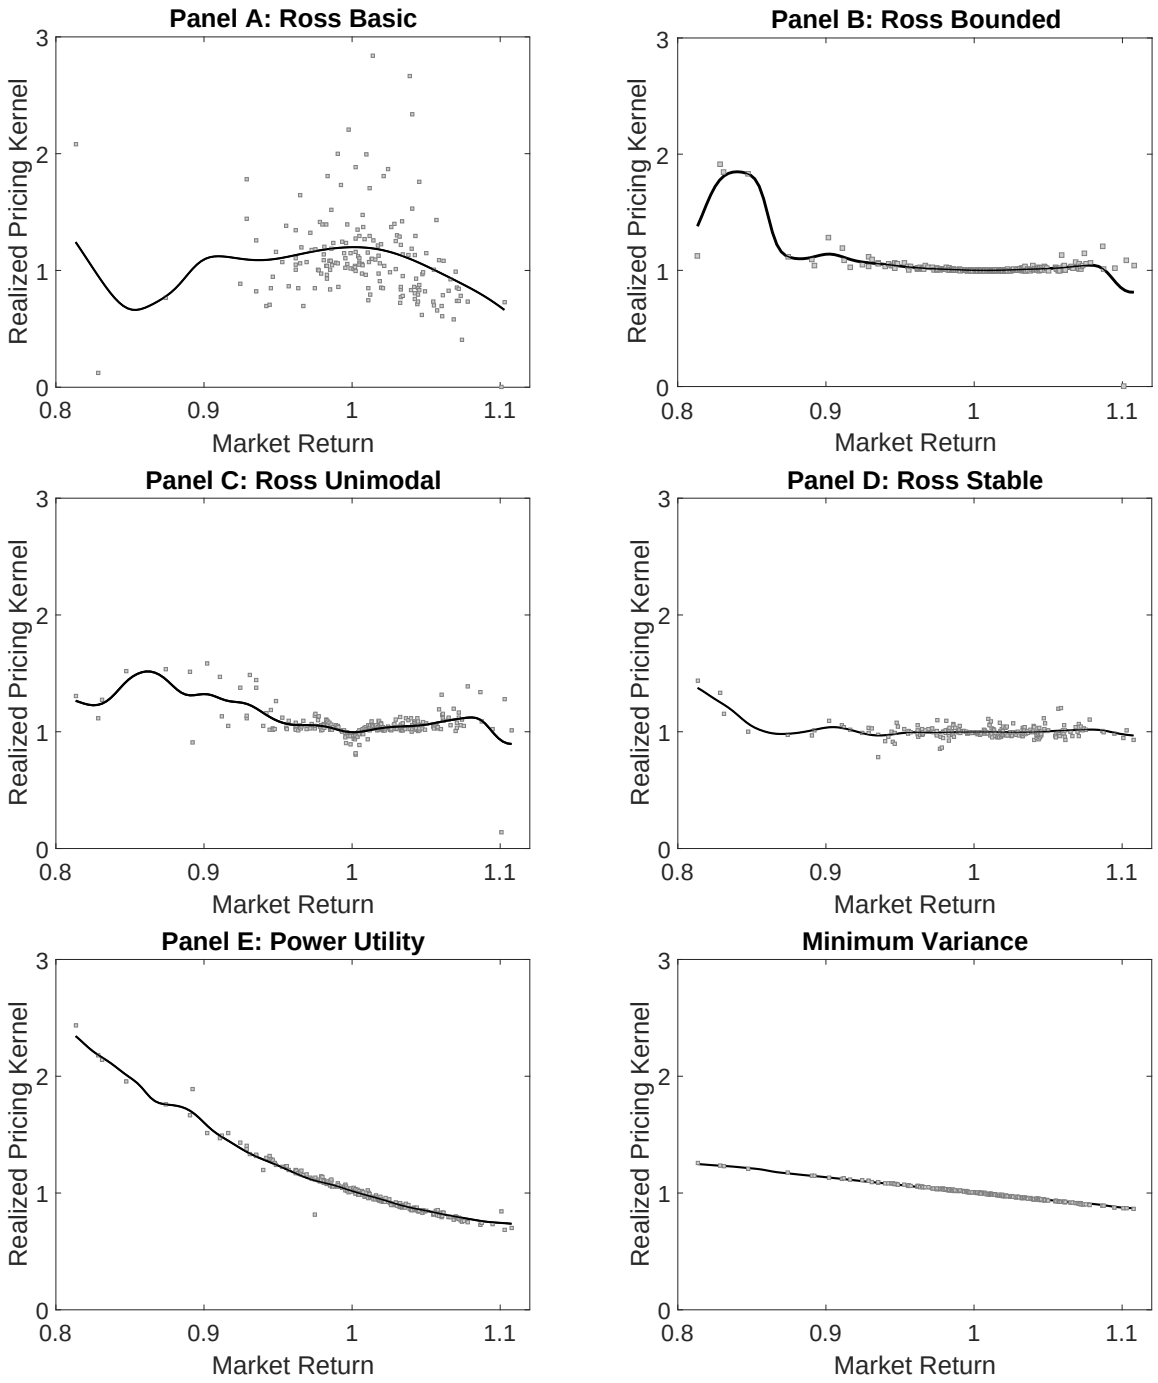

Figure 7. Time Series of Pricing Kernels. We depict a time series of realized pricing kernels (gray squares) for different recovery approaches. We further depict a Gaussian kernel regression line (black) through the realized pricing kernel values. Panel A shows realized pricing kernels for Ross Basic, Panel B for Ross Bounded, Panel C for Unimodal, Panel D for Ross Stable, Panel E for Power Utility with $\gamma = 4$, and Panel F for the Minimum Variance pricing kernel.

Utility generates a (more variable) pricing kernel, which broadly follows the shape of the Minimum Variance pricing kernel.

C. Decomposition of Pricing Kernels

Ross (2015) interprets his recovered probabilities $p_{i,j}$ as subjective beliefs of an investor whose preferences are reflected by the particular choice of pricing kernel $m_{i,j} = \delta u'_j/u'_i$. Borovicka et al. (2016) highlight that the recovered $p_{i,j}$ are not necessarily the true $p^{true}_{i,j}$. Only if the true pricing kernel, which translates true probabilities into state prices, were of exactly the functional form of Ross' assumption ($m_{i,j} = \delta u'_j/u'_i$), would the two probabilities coincide. Formally, Borovicka et al. (2016) decompose the state prices in the following way:

$$\pi_{i,j} = m_{i,j} \cdot p_{i,j} = m_{i,j} \cdot \frac{p_{i,j}}{p^{true}_{i,j}} \cdot p^{true}_{i,j} = m_{i,j} \cdot m^{perm}_{i,j} \cdot p^{true}_{i,j}, \tag{18}$$

with $m^{perm}_{i,j}$ being the ratio of recovered by true probabilities. The authors interpret the term $m^{perm}_{i,j}$ as a permanent component of the pricing kernel, which could be caused by permanent shocks to the macroeconomy. They label the Ross recovery pricing kernel m in contrast as the transitory pricing kernel component. They further note that Ross (2015) will only recover the true probabilities if $m^{perm}_{i,j} = 1$. As we use exactly this interpretation in our empirical work, we would like to know if the implicit assumption made by Ross (namely, that the permanent component of the pricing kernel is one) might contribute to the empirical failure of the recovery theorem. Empirically, Bakshi et al. (2016) suggest that $m^{perm}_{i,j}$ exhibits substantial variation and does not equal one. The authors further point out that the transitory pricing kernel component (which should equal our usual m from Ross recovery) can be found empirically, for date τ, as:

$$m^{\mathrm{T}}_{\tau} = \frac{1}{Rf_{\tau}} \cdot \frac{F_{\tau+,\infty}}{F_{\tau,\infty}}, \tag{19}$$

where Rf_{τ} is the 30-day risk-free return at date τ, $F_{\tau,\infty}$ is the price of a 30 year Treasury bond future at date τ, and $F_{\tau+,\infty}$ is the price of a 30 year Treasury bond future 30 calendar days after date τ.

To uncover if recovered pricing kernels empirically indeed equal the transitory component, we regress realized pricing kernels implied by a particular recovery approach on the empirical

transitory component, which we construct as in Equation 19:

$$m_\tau^{\text{Recovery}} = \beta_0 + \beta_1 m_\tau^{\text{T}} + \epsilon_\tau, \quad \tau = 1, ..., \mathcal{T}. \tag{20}$$

Table III. **Pricing Kernels and the Transitory Component.** We present the results from regressing the time series of realized pricing kernels for different recovery approaches (including Power Utility) on the time series of the empirical transitory component of the pricing kernel.

Recovery Approach	Intercept β_0	Slope β_1	adjusted R^2
Ross Basic $\pi_{i,j} > 0$	48.42	-40.33	-0.003
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.97***	0.05	-0.004
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.18***	-0.18	-0.003
Ross Stable Do not use transition state prices	0.86***	0.15	0.002
Power Utility with $\gamma = 4$	2.17***	-1.15**	0.020

* indicates significance at 10%, ** indicates significance at 5%, *** indicates significance at 1%.

Results for that regression are given in Table III. For Ross Basic, we again find a large negative yet insignificant intercept and a large positive yet insignificant slope coefficient. For Ross Bounded, Ross Unimodal, and Ross Stable we find positive and significant intercepts as well as insignificant slopes lower than one. The adjusted R^2 is very low for all these approaches with values below 0.002. Power Utility, having a positive intercept of 2.17, a negative slope of -1.15, and a low adjusted R^2 of 0.02, also exhibits very low similarity with the transitory component. However, that result does not discredit the Power Utility approach, as there is no theoretical reason why its pricing kernel should equal the transitory component. The story is different for the Ross recovery approaches. Their pricing kernels should theoretically match the transitory pricing kernel component, yet empirically do not.

D. *Time-Homogeneity of Transition State Prices and the Fitting of option prices*

From Figure 5 we learned that Ross Stable and Power Utility use the spot state prices directly, while Ross Basic, Ross Bounded, and Ross Unimodal fit the current transition state prices to those state prices, allowing for some discrepancy.[33] How large is that discrepancy? As a measure, we compare one-month observed implied volatilities σ^{obs} with one-month model implied volatilities $\hat{\sigma}$, where we use the one-month state prices of each approach to find model option prices and their associated model implied volatilities; see Appendix.D for details. For each date τ, we compute the root-mean-squared error between the two implied volatilities. Our final measure $MRMSE$ is the time series average of all the $RMSE_\tau$. Table IV shows the one-month $MRMSE$ for each recovery approach in columns 2 and 3.

We find the lowest one-month $MRMSE$ of 0.008 for the methods Ross Stable and Power Utility, which directly use the spot state prices and which do not use transition state prices. This is not surprising, as the error in these cases only stems from smoothing and interpolating the implied volatility surface. This low error also demonstrates the ability of our smoothing and interpolation methodologies to approximate observed implied volatilities reasonably well.

Among the transition state price versions, Ross Basic has only a slightly higher $MRMSE$ of 0.009. As we only require positivity of the transition state prices, the optimization can freely choose transition state prices to match the spot state price surface. This results in a good fit of one-month spot state prices but comes at the price of implying extreme monthly state-dependent risk-free rates.

For the other two transition state price approaches (Ross Bounded and Ross Unimodal), we find more than 15 times higher $MRMSEs$, 0.141 and 0.160. The transition state prices for those methods are almost useless for pricing, as the MRMSE is of the same magnitude as the implied volatility itself. Adding even mild economic constraints removes many degrees of freedom in the allocation of transition state prices. Therefore, choices for the transition state prices for the current state are more limited, increasing the $MRMSE$.

We further want to investigate how the assumption of time-homogeneous transition state prices influences our empirical results. Given time-homogeneity, we may multiply the one-month-maturity transition state price matrix Π t-times with itself to again generate transition state prices (and thus also spot state prices) with a maturity of t months. Empirically, we are concerned that the longer-dated transition state prices will not fit the option prices well, as

[33]Note that state prices and, thus, option prices are not defined for Historical Return Distribution.

Table IV. Accuracy of Transition State Prices. We present a measure for how closely the observed option prices are fitted by each of the five approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, and Power Utility with γ=4. Our measure is the time series average of the root-mean-squared errors of model versus observed implied volatilities, $MRMSE$. Columns 2 and 3 present the $MRMSE$ for a one-month maturity, where model implied volatilities are either based on transition state prices or spot state prices. Columns 4 and 5 present the $MRMSE$ for a twelve-month maturity, where model implied volatilities are either based on conflated transition state prices or spot state prices.

Recovery Approach	**One-month transition state prices** MRMSE	**One-month spot state prices** MRMSE	**12-month transition state prices** MRMSE	**12-month spot state prices** MRMSE
Ross Basic $\pi_{i,j} > 0$	0.009		0.050	
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.141		0.062	
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.160		0.065	
Ross Stable Do not use transition state prices		0.008		0.004
Power Utility with $\gamma = 4$		0.008		0.004
Historical Return Distribution	Not defined	Not defined	Not defined	Not defined

it is not obvious that longer-dated options are priced on the conflated state prices of shorter dated options.[34]

For the recovery approaches using transition prices (Ross Basic, Ross Bounded, and Ross Unimodal) and on every date τ, we conflate the transition state price matrix Π twelve times with itself to extract longer-dated spot state prices. Similarly to the one-month horizon, we compute the root-mean-squared differences of model and observed option implied volatilities. Last, we take the time series mean of the $RMSEs$ over all dates with available twelve-month maturity option data to arrive at the twelve-month $MRMSE$. For the approaches using the spot state prices directly (Ross Stable and Power Utility), we proceed analogously but use the twelve-month spot state prices instead of the transitions state prices. We report the results in columns 4 and 5 of Table IV.

We find the lowest twelve-month $MRSME$ with 0.004 for Ross Stable and Power Utility, being even lower that the one-month $MRMSE$ of those methods (0.008). The reason is that the longer-dated distributions cover a greater range of states and thus approximate the observed spot prices better. The twelve-month $MRMSE$ (0.062 and 0.065) for Ross Bounded and for Ross Unimodal are again some 15 times the values for Ross Stable and Power Utility. Most interestingly, while the transition state prices for Ross Basic were able to approximate one-month option prices well, the conflated 12-month transition state prices imply a large $MRMSE$ of 0.050.

We conclude that Ross Basic, Ross Bounded, and Ross Unimodal fail to produce transition spot state prices that approximate option prices well. Only Ross Basic prices the one-month options reasonably well but at the cost of implausibly shaped transition state prices and extreme risk-free interest rates.

E. Insights from Simulated Economies

We are concerned that small data errors in the option prices might cause the recovery theorem to fail. To check, we simulate economies, where a particular recovery approach holds true, and generate future realized returns by drawing from the recovered physical distribution. In that case and with a 5% significance level for our statistical tests, we should have a 95% non-rejection rate (i.e., future returns are compatible with the simulated recovery approach) and a 5% rejection rate.

[34]Again, to keep our discussion simple, we do not distinguish between risk-neutral probabilities and transition state prices because numerically, the small risk-free rate adjustment hardly makes any difference.

Next, we perturb the option prices. This will decrease the non-rejection rates as the perturbed recovery approaches generate physical distributions, which will deviate from the true physical distribution, from which we drew the future realized returns.

Our methodology is as follows. For each date τ in our sample, we assume that the true *risk-neutral* economy is represented by the observed implied volatilities $\sigma_{i(l),t(l)}$ with moneyness $i(l)$ and maturity $t(l)$ for observation $l = 1, ..., L$. At date τ, we further draw a one-month future return R_τ^{true} from the *physical* distributions p_τ^{true} that we recovered with a particular recovery approach. This gives us a time series of 223 drawn one-month future realized returns from an economy where the particular recovery approach holds.

On each date, we then add v times a standard normally distributed error ϵ to the observed implied volatilities $\sigma_{i(l),t(l)}$ and obtain *perturbed* implied volatilities $\hat{\sigma}_{i(l),t(l)}$ as:

$$\hat{\sigma}_{i(l),t(l)} = \sigma_{i(l),t(l)} + v \cdot \epsilon, \quad \epsilon \sim N(0,1). \tag{21}$$

We apply our standard algorithm to recover the physical distributions $\hat{p}_\tau$, which are now based on the *perturbed* implied volatilities $\hat{\sigma}_{i(l),t(l)}$. We then use a Knüppel test[35] to test the hypothesis that the simulated future returns R_τ^{true} are compatible with the recovered distributions $\hat{p}_\tau$. We draw 1000 realizations of future return time series, each consisting of 223 returns, and report the non-rejection rates for our hypothesis at the 5% significance level.

We use four different levels of perturbation, based on the multiplier v of the error term ϵ in Equation 21. We express v on each date τ as a multiple of the mean implied volatility bid/ask spread on that date.[36] Our four levels of perturbation are none, one-half, one, and twice the mean spread of a particular date.

Table V presents how often our hypothesis is accepted at the 5% significance level for different multipliers v of the error-term ϵ. Without any perturbation, the non-rejection rates are all very close to the theoretical value of 95%. Beyond that, we find that only Ross Basic turns out to be very sensitive to perturbations. Even though we simulate under the assumption that Ross Basic holds, we only have non-rejection rates of around 5% to 15% with perturbations instead of 95% without. Thus, the failure of Ross Basic could also be

[35] Using the Berkowitz test at the 5% level instead does not change our results.

[36] On each date, we compute the mean across all options of (ask implied volatility - bid implied volatility). The mean spread ranges from 0.007 to 0.061 through our 223 sample date with an average value of 0.016.

Table V. Stability of a Simulated Recovery Approaches. We generate 1000 realizations of an economy, where a particular recovery approach holds. For each realization, we draw 223 future realized returns. Next, we perturb option implied volatilities by adding v times a standard normally distributed error ϵ. Given v, we report the non-rejection rates for the hypothesis that the future realized returns were drawn from the perturbed physical distributions. We use a Knüppel test with a significance level of 5%.

Recovery Approach	$v = 0 \cdot$**spread** non-rej. rate	$v = \frac{1}{2} \cdot$**spread** non-rej. rate	$v = 1 \cdot$**spread** non-rej. rate	$v = 2 \cdot$**spread** non-rej. rate
Ross Basic	95.0 %	15.1 %	8.6 %	5.6 %
Ross Bounded	95.7 %	92.8%	88.8 %	61.4 %
Ross Unimodal	95.3 %	94.1 %	88.6 %	67.6 %
Ross Stable	95.0 %	95.6 %	93.4%	74.6%
Power Utility	96.4 %	91.7 %	88.4%	73.5 %
Hist. Return Distr.	n/a	n/a	n/a	n/a

driven by errors in the option implied volatilities.

Yet for all other approaches[37], the decay of the non-rejection rates is much slower as we increase our perturbations. Using the mean bid/ask spread at date τ as the multiplier v of our perturbation term ϵ, non-rejection rates are above 88.4%. Thus, Ross Bounded, Ross Unimodal, Ross Stable, and Power Utility should all be similarly sensitive to perturbation. However, only Power Utility cannot be rejected as a model of future realized returns, while we reject all Ross recovery approaches. We reason that the failure of Ross Bounded, Ross Unimodal, and Ross Stable is not likely to be driven by perturbations of the option prices.

We conclude that, while Ross Bounded, Ross Unimodal, and Ross Stable suffer from rather flat pricing kernels, the methods are less sensitive to perturbations of the option prices than Ross Basic. It seems that adding economic constraints results in models, which are less sensitive to perturbations. This finding also holds for Power Utility.

[37]The non-rejection rate cannot be computed for Historical Return Distribution.

VII. Robustness

To demonstrate the stability of our results, we implement two types of robustness checks. First, we investigate if variations in the state space significantly influence our empirical results. Second, we repeat our study but now exclude the periods related to (i) the Dot-Com bubble and (ii) the Financial Crisis from our original sample.

A. Variations in the State Space

The recovery theorem works on a finite state space and we are concerned that our methodological choices for the states might drive the results. We proceed to show that this is not the case when we use (i) log returns instead of straight returns or (ii) a 20% coarser state space than our usual fine state spaces with N=111 (120 for Ross Stable and Power Utility). If we (iii) use the original, non-overlapping state space with N=12 from the Ross (2015), then all recovery approaches and also Power Utility are being rejected. The fit to the option prices deteriorates markedly in the process. Note that the Historical Return Distribution is always unaffected by the choice of the state space and we thus always maintain our result that we cannot reject the Historical Return Distribution.

Log returns

We repeat our main empirical study but now define our state-space in log-returns. We construct our volatility surface by linearly interpolating the log-moneyness of the fine implied volatility surface. As before, we recover the physical distributions but based on the log state space. We then test if the one-month future realized log-returns are drawn from the recovered distributions. For the Historical Return Distribution we generate the empirical cumulative distribution on each sample date based on a five-year historical sample of one-month non-overlapping log-returns.

Table VI shows the resulting p-values for our three test statistics and can be readily compared to Table I. For Ross recovery, all p-values are even lower than in the main results except for Kolmogorov-Smirnov in case of Ross Basic, which increases from 0.000 to 0.025. We still reject our hypothesis that future realized returns are compatible with the recovered physical distributions for all recovery approaches. For Power Utility, we find lower p-values compared to the main results, but we still cannot reject our hypothesis at the 5% level with any of our three tests. The results for Historical Return Distribution remain virtually unchanged as the distribution is almost independent of the state space (except for some minute interpolation effect in case of the Berkowitz test). We do not report $MRMSE$

values for the log state space, as they are very close to the values of our main run. We conclude that our results are robust when changing to log returns.

Table VI. Tests of the recovered physical probabilities for log returns. We present our results if future realized log returns are drawn from physical probabilities generated by one of the six, log-scaled approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, Power Utility, and Historical Return Distribution. For each approach, we show the p-values from the Berkowitz, Knüppel, and Kolmogorov-Smirnov tests for uniformity of the percentiles of future realized log returns under the model physical cumulative distribution.

Recovery Approach	**Berkowitz** p-value	**Knüppel** p-value	**Kolmogorov-Smirnov** p-value
Ross Basic $\pi_{i,j} > 0$	0.001	0.013	0.025
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.000	0.000	0.000
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.000	0.000	0.000
Ross Stable Do not use transition state prices	0.000	0.002	0.000
Power Utility with $\gamma = 4$	0.285	0.159	0.082
Historical Return Distribution	0.290*	0.480	0.347

*Small discrepancies with the p-values of Table I are due to the linear interpolation of the cumulative density function of returns versus log returns.

Changing the dimension of the state-space

In a simulation study, Tran and Xia (2014) show that the recovered probabilities depend on the dimension of the state-space. In order to see if their concerns apply to our setting, we first reduce our overlapping fine state space by 20% and, second, we use a very coarse non-overlapping state space with just twelve states.

Our main study uses ten maturity steps per month which gives us 111 states for Ross

Basic, Ross Bounded, and Ross Unimodal and 120 states for Ross Stable and Power Utility. We now reduce the state space to eight maturity steps per month and end up with 89 states for Ross Basic, Ross Bounded, and Ross Unimodal and with 96 states for Ross Stable and Power Utility.

The results are shown in Table VII. Our results hardly change from the main run with the recovery approaches being strongly rejected at the 5% level (only the p-value for the Kolmogorov-Smirnov test for Ross Basic is 0.088) and the benchmark models not being rejected. We conclude that a reduction of our state dimension by 20% does not lead to significant changes of our test results.

Table VII. **Tests of the recovered physical probabilities for a reduced state-space size.** We present our results if future realized returns are drawn from physical probabilities generated by one of the six approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, Power Utility, and Historical Return Distribution. We reduce the original state-space-size by 20 %. For each approach, we show the p-values from the Berkowitz, Knüppel, and Kolmogorov-Smirnov tests for uniformity of the percentiles of future realized returns under the model physical cumulative distribution.

Recovery Approach	**Berkowitz** p-value	**Knüppel** p-value	**Kolmogorov-Smirnov** p-value
Ross Basic $\pi_{i,j} > 0$	0.026	0.035	0.088
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.001	0.002	0.001
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.001	0.000	0.002
Ross Stable Do not use transition state prices	0.000	0.000	0.006
Power Utility with $\gamma = 4$	0.643	0.256	0.340
Historical Return Distribution	0.294	0.480	0.347

Next, we use a non-overlapping state space with only twelve states (one maturity step per month), just as in the original paper of Ross (2015). For Ross Basic, Ross Bounded, and Ross Unimodal we thus use a twelve by twelve transition state price matrix. For Ross Stable, we solve Equation 14 to recover a one-month pricing kernel on a twelve state moneyness grid (one maturity step being one month).

Now we reject all four recovery approaches and even Power Utility with p-values close to zero. Only the results for Historical Return Distribution are not affected by the state space and we cannot reject our main hypothesis there. As the one-month spot state price approximation one twelve states is very coarse, the $MRMSE$ errors are also much higher with values beyond 0.100 for all five approaches.[38]

Choosing twelve equidistant states, which cover the range of the state prices even at the yearly horizon, means that the one-month state prices are only non-zero at a few states around the current state. Thus, we implement a second version based on a non-equidistant discretization with more of the twelve states being closer to the current state and fewer of the twelve states being further away (See Appendix.E for details). This version leads to a more accurate approximation of the one-month spot state prices at the cost of a worse approximation of higher maturity spot state prices. However, results are as disappointing as for the equidistant twelve state space. We conclude that small variations to the state space (20% smaller) do not affect the results, while going to a twelve state non-overlapping state space means that the spot state prices (and thus the option prices) can no longer be accurately fitted on such a coarse space.

B. Sub-Samples

For our main study, we use monthly data from January 1996 to August 2014 with a sample size of 223 dates. We repeat our main study but now consider two different sub-samples. In our first variation, we exclude the two years (February 1998 until February 2000) prior to the Dot-Com bubble burst in March 2000. Results are given in Table VIII. For our second variation, we follow Anand et al. (2013) and exclude the financial crisis period (July 2007 until March 2009). Results are given in Table IX.

In either case, results hardly change from our main results and we conclude that our study is robust if we exclude crucial periods such as the Dot-Com bubble or the financial crisis.

[38]Note that the error is not defined for Historical Return Distribution.

Table VIII. **Tests of the recovered physical probabilities excluding the period of the Dot-Com bubble.** We present our results if future realized returns are drawn from physical probabilities generated by one of our six approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, Power Utility, and Historical Return Distribution. We use our sample period from January 1996 until August 2014 but exclude months that are associated with the Dot-Com bubble (February 1998 until February 2000). For each approach, we show the p-values from the Berkowitz, Knüppel, and Kolmogorov-Smirnov tests for uniformity of the percentiles of future realized returns under the model physical cumulative distribution.

Recovery Approach	**Berkowitz** p-value	**Knüppel** p-value	**Kolmogorov-Smirnov** p-value
Ross Basic $\pi_{i,j} > 0$	0.016	0.039	0.000
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.007	0.003	0.014
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.001	0.000	0.027
Ross Stable Do not use transition state prices	0.015	0.024	0.007
Power Utility with $\gamma = 4$	0.552	0.378	0.561
Historical Return Distribution	0.745	0.807	0.575

Table IX. **Tests of the recovered physical probabilities excluding the period of the financial crisis.** We present our results if future realized returns are drawn from physical probabilities generated by one of our six approaches: Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, Power Utility, and Historical Return Distribution. We use our sample period from January 1996 until August 2014 but exclude months that are associated with the financial crisis (July 2007 until March 2009). For each approach, we show the p-values from the Berkowitz, Knüppel, and Kolmogorov-Smirnov tests for uniformity of the percentiles of future realized returns under the model physical cumulative distribution.

Recovery Approach	**Berkowitz** p-value	**Knüppel** p-value	**Kolmogorov-Smirnov** p-value
Ross Basic $\pi_{i,j} > 0$	0.047	0.021	0.005
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.001	0.003	0.004
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.000	0.000	0.012
Ross Stable Do not use transition state prices	0.002	0.010	0.002
Power Utility with $\gamma = 4$	0.615	0.359	0.198
Historical Return Distribution	0.752	0.700	0.394

VIII. Conclusion

We implement and test several recovery approaches in the framework of Ross (2015). We further present a variant of Ross recovery without explicitly estimating the transition state prices, which are numerically hard to determine. We find that future realized S&P 500 returns are incompatible with the recovered physical probabilities. On the other hand, two simple benchmark approaches work much better. For one, we employ the empirical distribution of non-overlapping monthly S&P 500 returns during the past five years. Alternatively, we use the pricing kernel of a power utility function with risk aversion parameter $\gamma = 4$. Both simple benchmark models cannot be rejected in the data.

We further analyze why the recovery theorem fails. We find that the most basic approach, requiring minimal assumptions, delivers unstable transition state prices. Backed by a simulation study, we argue that it is very sensitive to small variations in the options data to which the model is fitted. Further, the extreme risk-free rates implied by this basic approach are economically implausible.

Alternative implementations of Ross recovery, with added economically reasonable constraints on the transition state prices, are much more stable than the basic approach but are not able to generate pricing kernels that are sufficiently away from risk-neutrality to being compatible with future realized index returns.

Robustness tests confirm that all Ross recovery approaches are incompatible with future realized S&P 500 index returns. Further, the assumption of time-homogeneity of the transition state prices leads to poorly fitted option prices for all Ross recovery approaches. Then, we check on an implicit assumption of Ross, which sets the permanent pricing kernel component to one and leads to the recovery of only the transitory pricing kernel component. Comparing the recovered transitory component to an empirically estimated transitory component as in Bakshi et al. (2016), we find that the two approaches do not coincide.

REFERENCES

Ait-Sahalia, Y. and Lo, A. W. (2000). Nonparametric Risk Management and Implied Risk Aversion. *Journal of Econometrics*, 94(1-2):9–51.

Alvarez, F. and Jermann, U. J. (2005). Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth. *Econometrica*, 73(6):1977–2016.

Anand, A., Irvine, P. J., Puckett, A., and Venkataraman, K. (2013). Institutional Trading and Stock Resiliency: Evidence from the 2007-2009 Financial Crisis. *Journal of Financial Economics*, 108(3):773–797.

Andrews, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, 59(3):817–858.

Audrino, F., Huitema, R., and Ludwig, M. (2015). An Empirical Analysis of the Ross Recovery Theorem. *University of St. Gallen, University of Zurich, University of Zurich. Working Paper, Available at SSRN*.

Bakshi, G., Chabi-Yo, F., and Gao, X. (2016). A Recovery That We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem. *University of Maryland, University of Massachusetts Amherst, University of Maryland. Working Paper, Available at SSRN*.

Berkowitz, J. (2001). Testing Density Forecasts, With Applications to Risk Management. *Journal of Business and Economic Statistics*, 19(4):465–474.

Bliss, R. and Panigirtzoglou, N. (2004). Option-Implied Risk Aversion. *Journal of Finance*, 59(1):407–446.

Borovicka, J., Hansen, L. P., and Scheinkman, J. A. (2016). Misspecified Recovery. *Journal of Finance*, 71(6):2493–2544.

Bowman, A. W. and Azzalini, A. (1997). Applied Smoothing Techniques for Data Analysis. *New York: Oxford University Press Inc.*

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of State-Contingent Claims implicit in Option Prices. *Journal of Business*, 51(4):621–651.

Carr, P. and Yu, J. (2012). Risk, Return, and Ross Recovery. *Journal of Derivatives*, 20(1):38–59.

Christensen, T. M. (2017). Nonparametric Stochastic Discount Factor Decomposition. *New York University. Working Paper, Available at https://arxiv.org/pdf/1412.4428.pdf.*

Cochrane, J. H. (2000). Asset Pricing. *Princeton University Press, Princeton, USA, 1 edn.*

Cuesdeanu, H. and Jackwerth, J. C. (2016). The Pricing Kernel Puzzle in Forward Looking Data. *University of Konstanz. Working Paper, Available at SSRN.*

Dubynskiy, S. and Goldstein, R. S. (2013). Recovering Drifts and Preference Parameters from Financial Derivatives. *University of Minnesota. Working Paper, Available at SSRN.*

Ghosh, A., Julliard, C., and P, T. A. (2017). What Is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models. *Review of Financial Studies*, 30(2):442–504.

Hansen, L. P. and Scheinkman, Jose, A. (2009). Long Term Risk: An Operator Approach. *Econometrica*, 77(1):173–234.

Jackwerth, J. C. (2000). Recovering Risk Aversion from Option Prices and Realized Returns. *Review of Financial Studies*, 13(2):433–451.

Jackwerth, J. C. (2004). Option-Implied Risk-Neutral Distributions and Risk Aversion. *Research Foundation of AIMR, CFA Institute.*

Jensen, C. S., Lando, D., and Pedersen, Lasse, H. (2017). Generalized Recovery. *Copenhagen Business School. Working Paper, Available at SSRN.*

Knüppel, M. (2015). Evaluating the Calibration of Multi-Step-Ahead Density Forecasts Using Raw Moments. *Journal of Business and Economic Statistics*, 33(2):270–281.

Martin, I. and Ross, S. (2013). The Long Bond. *Stanford GSB, MIT Sloan. Working Paper.*

Massacci, F., Williams, J., and Zhang, Y. (2016). Empirical Recovery: Hansen-Scheinkman Factorization and Ross Recovery from High Frequency Option Prices. *University of Trento, Durham Business School. Working Paper, Available at SSRN.*

Qin, L. and Linetsky, V. (2016). Positive Eigenfunctions of Markovian Pricing Operators: Hansen-Scheinkman Factorization, Ross Recovery and Long-Term Pricing. *Operations Research*, 64(1):99–117.

Qin, L., Linetsky, V., and Nie, Y. (2016). Long Forward Probabilities, Recovery and the Term Structure of Bond Risk Premiums. *Northwestern University. Working Paper.*

Rosenberg, J. V. and Engle, R. F. (2002). Empirical Pricing Kernels. *Journal of Financial Economics*, 64(3):341–372.

Ross, S. (2015). The Recovery Theorem. *Journal of Finance*, 70(2):615–648.

Schneider, P. and Trojani, F. (2016). (Almost) Model-Free Recovery. *University of Lugano, University of Geneva. Working Paper, Available at SSRN.*

Tran, N.-K. and Xia, S. (2014). Specifed Recovery. *Washington University in St. Louis.*

Walden, J. (2016). Recovery with Unbounded Diffusion Processes. *Review of Finance, forthcoming.*

Appendices

A. *Smoothing the Implied Volatility Surface*

Analogous to Jackwerth (2004), we minimize the sum of squared local total second derivatives of implied volatilities (insuring smoothness of the volatility surface) plus the sum of squared deviations of the observed from model implied volatilities (insuring option fit) by using a trade-off parameter λ. The optimization problem is then defined as:

$$\begin{aligned} \min_{\sigma_{i,t}} \quad & \frac{1}{TN} \cdot \sum_{t=1}^{T} \sum_{i \in I} (\sigma''_{i,t})^2 \cdot t \quad + \quad \lambda \cdot \frac{1}{L} \cdot \sum_{l=1}^{L} \left(\sigma_{i(l),t(l)} - \sigma^{\text{obs}}_{i(l),t(l)} \right)^2 \\ & \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad s.t. \\ & \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \sigma_{i,t} \;\geq\; 0, \end{aligned} \tag{22}$$

where $\sigma''_{i,t}$ is the local second derivative (to be defined momentarily) and where t compensates for the lower curvature at higher maturities. $\sigma^{\text{obs}}_{i(l),t(l)}$ is the $l-th$ observed implied volatility with a total number of L observations. We start with a high trade-off parameter λ and iteratively increase the smoothness of the volatility surface by reducing λ, thus reducing the fit of observed implied volatility, until we obtain a smooth and positive state price surface.

We define the local second derivatives $\sigma''_{i,t}$ of implied volatilities as:

$$\begin{aligned} \sigma''_{i,t} \;=\; & \frac{\sigma_{i+1,t} - 2\sigma_{i,t} + \sigma_{i-1,t}}{(\Delta_i)^2} + \frac{\sigma_{i,t+1} - 2\sigma_{i,t} \;+\; \sigma_{i,t-1}}{(\Delta_t)^2} \\ & + \frac{\sigma_{i+1,t+1} - \sigma_{i+1,t-1} - \sigma_{i-1,t+1} + \sigma_{i-1,t-1}}{(4\Delta_i \Delta_t)}, \end{aligned} \tag{23}$$

where the first and second terms are approximations of the partial second derivative with respect to moneyness and maturity. The third term approximates the cross derivative. We evaluate the local second derivatives $\sigma''_{i,t}$ on a fine equidistant grid at time t and moneyness indexed by $i \in I$, where $I = \{-N^{fine}_{low}, ..., N^{fine}_{high}\}$ with N states. Step sizes are Δt for time and Δi for moneyness. We impose the following boundary conditions for all i and t:

$$\begin{aligned} & \sigma_{i,0} = \sigma_{i,1}, \quad \sigma_{i,T+1} = \sigma_{i,T} \\ \text{and} \quad & \sigma_{i-1,t} = \sigma_{i,t}, \text{ for } i = -N^{fine}_{low}, \quad \sigma_{i+1,t} = \sigma_{i,t}, \text{ for } i = N^{fine}_{high} \\ \text{and} \quad & \sigma_{i-1,T+1} = \sigma_{i,T}, \text{ for } i = -N^{fine}_{low}, \quad \sigma_{i+1,0} = \sigma_{i,1}, \text{ for } i = N^{fine}_{high} \\ \text{and} \quad & \sigma_{i-1,0} = \sigma_{i,1}, \text{ for } i = -N^{fine}_{low}, \quad \sigma_{i+1,T+1} = \sigma_{i,T} \text{ for } i = N^{fine}_{high}. \end{aligned} \tag{24}$$

B. *Recovery without Using Transition State Prices*

We start with Equation 12 and show how it can be used to achieve recovery without explicitly deriving transition state prices. We restate the equation in greater detail:

$$\begin{pmatrix} \pi^t_{-N_{low},-N_{low}} & \pi^t_{-N_{low},-N_{low}+1} & \cdots & \pi^t_{-N_{low},N_{high}-1} & \pi^t_{-N_{low},N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^t_{0,-N_{low}} & \pi^t_{0,-N_{low}+1} & \cdots & \pi^t_{0,N_{high}-1} & \pi^t_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^t_{N_{high},-N_{low}} & \pi^t_{N_{high},-N_{low}+1} & \cdots & \pi^t_{N_{high},N_{high}-1} & \pi^t_{N_{high},N_{high}} \end{pmatrix} \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix} = \delta^t \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix}, \tag{25}$$

where $\pi^t_{i,j}$ is the transition state price of moving from state i to state j over t transition periods. Note that the current row in Π^t (indexed by $i = 0$) represents the transition state prices with maturity equal to t transition periods:

$$\begin{pmatrix} \pi^t_{0,-N_{low}} & \pi^t_{0,-N_{low}+1} & \cdots & \pi^t_{0,N_{high}-1} & \pi^t_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix} = \delta^t \cdot z_0. \tag{26}$$

After dividing both sides by z_0, we use Equation 26 for every $t = 1, ..., T$ and stack the resulting equations to obtain the following system of equations:

$$\begin{pmatrix} \pi^1_{0,-N_{low}} & \pi^1_{0,-N_{low}+1} & \cdots & \pi^1_{0,N_{high}-1} & \pi^1_{0,N_{high}} \\ \pi^2_{0,-N_{low}} & \pi^2_{0,-N_{low}+1} & \cdots & \pi^2_{0,N_{high}-1} & \pi^2_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{T-1}_{0,-N_{low}} & \pi^{T-1}_{0,-N_{low}+1} & \cdots & \pi^{T-1}_{0,N_{high}-1} & \pi^{T-1}_{0,N_{high}} \\ \pi^T_{0,-N_{low}} & \pi^T_{0,-N_{low}+1} & \cdots & \pi^T_{0,N_{high}-1} & \pi^T_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} \frac{z_{-N_{low}}}{z_0} \\ \vdots \\ \frac{z_{-1}}{z_0} \\ 1 \\ \frac{z_1}{z_0} \\ \vdots \\ \frac{z_{N_{high}}}{z_0} \end{pmatrix} = \begin{pmatrix} \delta \\ \delta^2 \\ \vdots \\ \delta^{T-1} \\ \delta^T \end{pmatrix}. \tag{27}$$

To solve this system of equations, we note that all transition state prices starting at the current state but with different maturities can be replaced by the spot state prices of the

same maturity, which can be directly obtained from the spot state price surface. The new system of equations has N unknowns $\left(\frac{z_j}{z_0} \text{ with } j \in I \text{ and } \delta\right)$ and as many equations.[39] As δ is the utility discount factor, we require it to be larger than zero and smaller than one. We force the ratios $\frac{z_j}{z_0}$ to be non-negative, as these ratios are directly linked to the pricing kernel. Then, we solve the equation system by means of least squares.

We now prove equation 15 by restarting our discussion with equation 12 from above, which we repeat here for convenience:

$$\Pi^t z = \delta^t z \text{ with } t = 1, ..., T. \tag{28}$$

Solving the Eigenvalue problem, the pricing kernel for state j with a maturity of t transition periods can be found as:

$$m^t_{0,j} = \delta^t \frac{z_0}{z_j}, \quad j \in I. \tag{29}$$

The pricing kernels for different maturities of t transition periods thus only differ from the fraction $\frac{z_0}{z_j}$ by a factor involving t and the discount factor δ. We can use this property and link pricing kernels with different maturities in the following way:

$$m^t_{0,j} = \delta^t \frac{z_0}{z_j} = \delta^{t-1}\delta \frac{z_0}{z_j} = \delta^{t-1} m_{0,j}, \quad j \in I, \tag{30}$$

where $m_{0,j}$ represents the pricing kernel value for state j with a maturity of one transition period.

C. Test Statistics

Berkowitz test

Formally, we write the transformation of returns R_τ into x_τ as:

$$x_\tau = \hat{P}_\tau(R_\tau) = \int_{-\infty}^{R_\tau} \hat{p}_\tau(v) dv, \tag{31}$$

where $x_\tau \sim i.i.d.\ U(0,1)$. The Berkowitz (2001) test jointly tests uniformity and the i.i.d. property of x_τ. For this test, the series x_τ is transformed by applying the inverse standard

[39] Note that I has N elements, but there are only $N-1$ fractions as $\frac{z_0}{z_0} = 1$.

normal cumulative density function Φ to x_τ:

$$z_\tau = \Phi^{-1}(x_\tau) = \Phi^{-1}\left(\int_{-\infty}^{R_\tau} \hat{p}_\tau(v)dv\right).$$

Under the hypothesis $\hat{p}_\tau = p_\tau$, z_τ is distributed standard normally, which suggests the following AR(1) model:

$$z_\tau - \mu = \rho(z_{\tau-1} - \mu) + \epsilon_\tau, \tag{32}$$

where the null requires $\mu = 0$, $\text{Var}(\epsilon_\tau) = 1$, and $\rho = 0$. Berkowitz then applies the following likelihood ratio test:

$$LR_3 = -2(LL(0, 1, 0) - LL(\hat{\mu}, \hat{\sigma}, \hat{\rho})),$$

where LL characterizes the log likelihood of Equation 32.

Knüppel test

The Knüppel (2015) test first scales the series x_τ to $y_\tau = \sqrt{12}\,(x_\tau - 0.5)$. In order to test x_τ for standard uniformity, the series y_τ is tested for scaled uniformity with mean=0 and variance=1. The test then compares the first S moments of the series y_τ to the respective theoretical moments in the following GMM-type procedure with the test statistic α_S:

$$\alpha_S = \mathcal{T} \cdot D_S^\intercal \Omega_S^{-1} D_S, \tag{33}$$

where D_S is a vector that consists of the differences between the sample moments $\frac{1}{\mathcal{T}} \sum_{\tau=1}^{\mathcal{T}} y_\tau^s$ and the theoretical moments m_s for $s = 1, ..., S$. Ω_S is a consistent covariance matrix estimator of all S respective moment differences. We follow Knüppel (2015) and set all elements of Ω_S which represent covariances between odd and even moment differences to zero and apply the test by considering the first four moments (S=4). We account for serial correlation of x_τ by estimating a Newey-West covariance matrix with automated lag length as proposed by Andrews (1991).

Kolmogorov-Smirnov test

The Kolmogorov-Smirnov test looks at the maximum distance between the empirical $\hat{P}$ and theoretical P cumulative density function and uses the following test statistic:

$$KS = \sup_{v} |P(v) - \hat{P}(v)|. \tag{34}$$

D. In-Sample Pricing Error for Recovery

To obtain the mean-root-mean-squared-error ($MRMSE$) for implied volatilities for a specific recovery approach, we use the one-month spot state prices $\hat{\pi}$ that are implied by the recovery approach. For sample date τ and for each moneyness level K_l where we observe a one-month option price, we compute one-month implied option prices $\hat{C}$ that are based on state prices $\hat{\pi}$ by numerical integration:

$$\hat{C}(K_l) = \int_0^\infty \hat{\pi}(S) \cdot \max(S - K_l, 0) dS. \tag{35}$$

We then transform the option prices $\hat{C}(K_l)$ into implied volatilities $\hat{\sigma}(K_l)$. For each date τ, we compute the root-mean-squared error $RMSE_\tau$ between one-month observed implied volatilities and model implied volatilities $\hat{\sigma}(K_l)$. Our final measure $MRMSE$ is the time series average of all the $RMSE_\tau$. Analogous, we determine the one-year $MRMSE$ with one-year spot state prices $\hat{\pi}$ that are based on a particular recovery approach.

E. Non-Equidistant Interpolated Twelve-by-Twelve Spot State Price Space

Our usual twelve-by-twelve state space is equidistant in maturity and moneyness. For a non-equidistant state space in the moneyness dimension, we initially need 14 moneyness level, which will be reduced by the Breeden-Litzenberger approach to the final twelve moneyness levels. We pick the first six moneyness levels at: the beginning $N1$ and the end $N12$ of the moneyness range needed to cover the tails of *one-year* spot state prices, the beginning $N3$ and the end $N10$ of the moneyness range needed to cover the tails of *one-month* spot state prices, a moneyness of zero $N0$ where we assign the same implied volatility value as at $N1$, and a moneyness of three $N13$ where we assign the same implied volatility value as at $N12$.

We optimize for the location of the eight remaining moneyness levels in a way that the

least squares distance between the interpolated volatility surface and the volatility surface on the fine grid is minimized, while we require $N2$ to lie in between $N1$ and $N3$, $N4$ to $N9$ (with one of them being the current state with moneyness of one) to lie in between $N3$ and $N10$, and $N11$ to lie in between $N10$ and $N12$.