Does the Ross Recovery Theorem work Empirically?

Jens Carsten Jackwerth* Marco Menner†

June 24, 2016

Abstract

Starting with the fundamental relationship that state prices are the product of physical probabilities and the pricing kernel, Ross (2015) shows that, given severe assumptions, knowing state prices suffices for backing out physical probabilities and the pricing kernel at the same time. We empirically investigate the performance of such recovery based on the S&P 500 index by testing if future, realized returns are indeed compatible with the recovered physical distributions. Statistical tests are strongly rejected in our sample and we document a number of empirical problems.

Keywords: Ross recovery, Berkowitz, pricing kernel, risk-neutral density, 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. We thank 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]

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

Present research picks 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 Cuesdeanu and Jackwerth (2016), starts out with risk-neutral and physical probabilities in order to find the empirical pricing kernel. 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. Our contribution is to empirically analyze his theorem 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 detail the underlying assumptions of Ross (2015) some more.

First, the Ross recovery theorem 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 readily be found from option prices, see e.g. Jackwerth (2004). Yet Ross recovery also requires the transition state prices emanating from alternative, hypothetical states of the world.[2] There is, however, no market with actively traded financial

[1] See e.g. Cochrane (2000), pp. 50. A state (Arrow-Debreu) price represents the amount of dollars 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 a representative agent.

[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

products which allow to back out (non-spot) transition state prices $\pi_{i,j}$ directly. We therefore use information on spot state prices with N different maturities to obtain transition state prices linking N states.[3] This requires us to first estimate a rich 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 methods. We also introduce a version that allows to apply the recovery theorem directly to spot state prices $\pi_{0,j}$.

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.[4]

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 $p_{0,j}$. In each month t, 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_t for each month t that lies in between 0 and 1. If our hypothesis holds, the set $\{x_t\}$ is uniformly distributed. Following Berkowitz (2001)[5], we can strongly reject our hypothesis. Ross (2015) does not recover physical spot distributions of returns, which are consistent with future, realized S&P 500 returns. We further find that the recovery theorem does not produce reasonable pricing kernels as they are riddled with local minima and maxima. In contrast, we cannot reject the hypotheses that future 500 returns are drawn from

one period later.

[3]For implementation details and variations, see below.

[4]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)

[5]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.

simple physical distributions based on a power pricing kernel or the five-year historical return distribution.

The empirical problems of Ross recovery stem from three sources. For one, the numerical methods for obtaining transition state prices from (noisy) option prices add estimation error to the state prices. Second, the assumption of time-homogeneous transition state prices might not hold and different periods might require different transition state prices. Third, the strong assumption concerning the functional form of the pricing kernel is rather limiting. 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]

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 while staying close to the original methods of Ross (2015). They differ from Ross (2015) and us in that they use a neural network to interpolate the state price surface instead of smoothing techniques. Surprisingly, their recovered pricing kernels tend to be rather smooth and u-shaped, as opposed to our wavy pricing kernels. The subtle reason is a quadratic penalty term which they use to force all transition state prices to zero, but which, empirically, 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 distance to the current state, which in turn mechanically leads to u-shaped pricing kernels. 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.

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

[6]Moreover, under Ross recovery, pricing kernels are scaled differently according to each initial state, while CRRA pricing kernels (using returns as states) and CARA pricing kernels (using prices as states) do not depend on the initial state.

with the spot state prices. Jensen et al. (2016) suggest the same method and further restrictions of the pricing kernel. Their focus is then on using the mean of the recovered physical probabilities to predict S&P 500 return with a small yet significant R^2 of 1.28%. While they also apply a Berkowitz test and reject their particular model, they, in contrast to our work, do not cover the original formulation of Ross (2015).

Informing our empirical work is a simulation study by Tran and Xia (2014), who show that the recovered physical probabilities vary substantially for different state space dimensions. We check for this prediction in our robustness tests but it empirically does not matter much for our paper. Related to Ross recovery, Hansen and Scheinkman (2009) already showed theoretically that the Perron-Frobenius theorem can be used to identify a probability measure $\tilde{p}$ from a risk-neutral process, where $\tilde{p}$ contains information on long-term risk. Yet they argue that this measure typically differs from the true physical physical measure p. Borovicka et al. (2015) build on that approach and show that Ross, who recovers that probability measure $\tilde{p}$, ignores a non-trivial martingale component that links the pricing kernel associated with $\tilde{p}$ and the pricing kernel associated with the true physical probability measure p. Bakshi et al. (2015) support this result with the empirical finding that this martingale component indeed exists. Yet, these papers point out that the recovered probability measure might be useful in analyzing long-term bond prices. Turning this argument around, Martin and Ross (2013) show that the recovered pricing kernel can be inferred from the time series behavior of the long bond and information on the long end of the yield curve.

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 (2014) 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. Qin and Linetsky (2015) extend the recovery theorem to continuous-time with state variables that follow a Borel-right process. This class contains relevant models in finance such as multi-dimensional diffusion models and affine jump

models for both bounded and unbounded state spaces. Dubynskiy and Goldstein (2013) explain that bounding the state variable process is the key mechanism that makes the recovery of drifts and risk aversion possible. They argue that such boundary conditions have a significant impact on the recovery result and should only be applied if they can be justified economically.

Further work on the recovery of subjective probabilities and the pricing kernel is done by Schneider and Trojani (2015). They express the pricing kernel as a polynomial function of both risk-neutral and subjective moments. They minimize its expected variance and require mild economical assumptions such as a negative divergence premium and bounded conditional subjective moments to achieve recovery.

The remainder of the paper proceeds as follows. Section 2 explains the Ross recovery theorem. In Section 3, 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, which we use as the main test in our empirical approach, as well as additional statistical tests which we apply for robustness. Section 4 describes our data set. In section 5, we present the empirical results of our study. Section 6 provides several robustness checks. Section 7 contains concluding remarks.

II. The Ross Recovery Theorem

An application of the recovery theorem requires the following assumptions: **(i)** Transition state prices $\pi_{i,j}$ need to be strictly positive, **(ii)** the transition state prices follow a time homogeneous process, which means that they are independent of calender time, 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 (marginal utility) values u'_j and u'_i.

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 and all the u'_i from the transition state prices $\pi_{i,j}$ and thus can uniquely solve for the physical transition probabilities.

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 $N = N_{low} + N_{high} + 1$. States lower than the current state are indexed with a negative i, while states higher than the current state are indexed with a positive i. The ending transition

state j is drawn from the same set I. Analogous to the two state example, we solve the following N-dimensional eigenvalue problem:

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

Knowing z_i and δ, we are then able to recover the physical transition probabilities $p_{i,j}$. The $0-th$ row of Π represents the spot state prices $\pi_{0,j}$, which coincide for the first period with the transition state prices.

III. Methodology

We now explain how we estimate a smooth implied volatility surface based on observed option prices. We then transform this smooth implied volatility surface into a smooth call option price surface. From this call option price surface, we obtain spot state prices $\pi_{0,j}^{(t)}$ for different maturities t, which we then use to back out transition state prices $\pi_{i,j}$.

A. Smoothing the Implied Volatility Surface

We use European put- and call options on the S&P 500 for our analysis and transform observed bid- and ask option prices into implied volatility bid- and ask quotes. Quotes are only available for some specific states and maturities, but we are interested in a rich and smooth implied volatility surface. Jackwerth (2004) introduces the fast and stable method to find smooth implied volatilities σ_i on a fine grid of states i for a fixed maturity. He 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 a 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 the volatility smile) plus the sum of

squared violations of the observed bid and ask implied volatilities (insuring the fit to the options data by penalizing model values outside the bid/ask bounds) using a trade-off parameter λ. The optimization problem is then defined as:

$$\begin{aligned} \min_{\sigma_{i,t},\ \delta_l^{\text{bid}},\ \delta_l^{\text{ask}}} \quad & \frac{1}{TN} \cdot \sum_{t=1}^{T} \sum_{i=1}^{N} (\sigma''_{i,t})^2 \quad + \quad \lambda \cdot \frac{1}{L} \cdot \sum_{l=1}^{L} \left(\delta_l^{\text{bid}} + \delta_l^{\text{ask}} \right)^2 \\ & s.t. \quad \delta_l^{ask} \ \geq \ \sigma_{i_l,t_l} - \sigma_{i_l,t_l}^{ask} \\ & \qquad \delta_l^{bid} \ \geq \ \sigma_{m_l,t_l}^{bid} - \sigma_{i_l,t_l} \\ & \qquad \sigma_{i,t},\ \delta_l^{\text{bid}},\ \delta_l^{\text{ask}},\ \geq \ 0 \end{aligned} \tag{7}$$

where σ_{i_l,t_l}^{bid} and σ_{i_l,t_l}^{ask} are the $l-th$ observed implied volatility bounds with a total number of L observations. δ_l^{bid} represents the distance of the implied volatility, if it is below the bid bound and δ_l^{ask} represents the distance of the implied volatility, if it is above the ask bound. One advantage of that method is that implied volatility can move flexible within bid- and ask bounds. We transform the resulting implied volatilities into call option prices and, by means of Breeden and Litzenberger (1978) and for each maturity t, take the numerical second derivative of call prices to obtain spot state prices. We insure positivity and smoothness of the spot state price surface by reducing the trade-off parameter λ in (7) until our imposed conditions for smoothness and positivity of the state price surface are fulfilled. Appendix A provides details for these conditions and the implementation.

We discretize the maturity at a step size of 3.5 days and let it move within a range of 6.5 days and 365 days. We round observed option expiry dates to the respective next maturity that lies on that grid. For the state space dimension, we discretize option strike prices with a step size of \$5. We then convert strike prices into moneyness by normalizing them with the current level of the S&P 500 index. This discretization enables to use all available information on option prices within the defined maturity range. We set sufficient upper and lower bounds for the moneyness such that all possible states with a non-zero probability of occurrence are covered.

Figure 1 shows the smoothed implied volatility surface on February 17, 2010, while Figure 2 illustrates spot state prices based on the smoothed implied volatility surface.

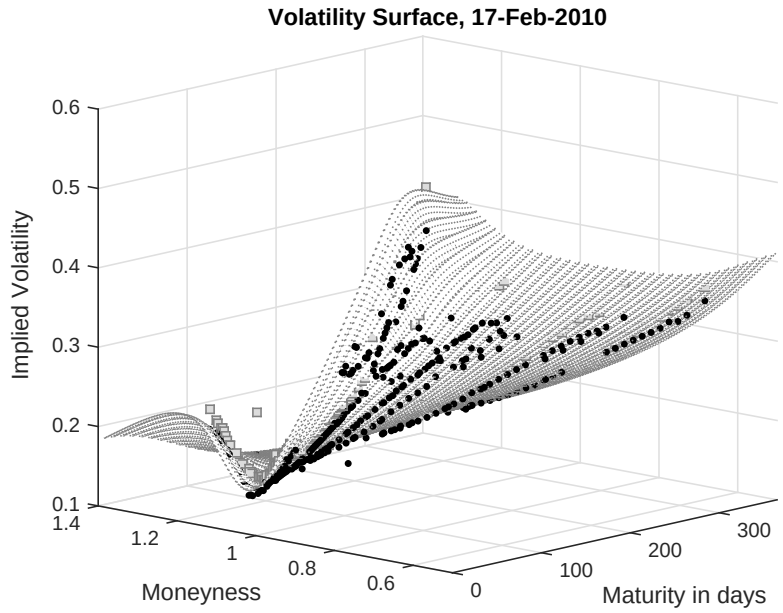

Figure 1. Implied volatility surface obtained by our smoothing method on February 17, 2010. Bid- and ask observed implied volatilities are marked as black dots and gray squares.

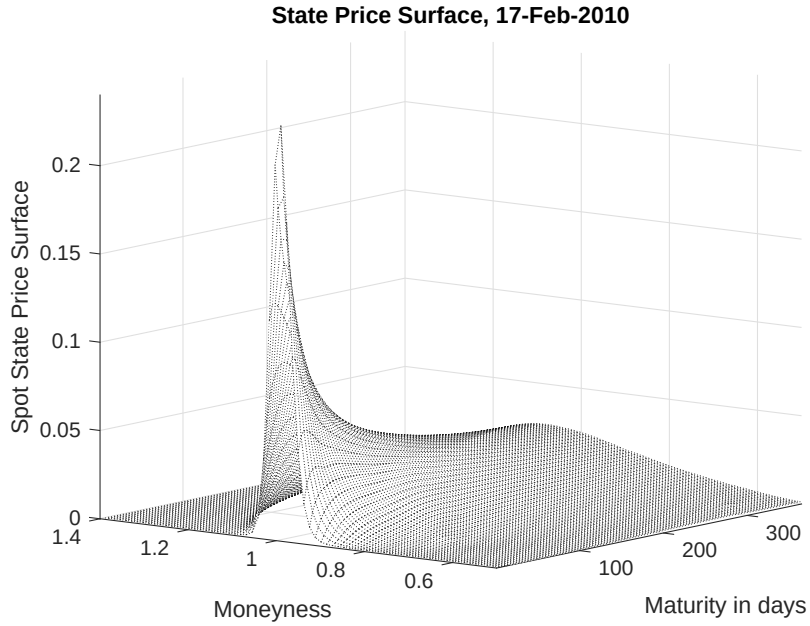

Figure 2. Spot state prices implied by the smoothed implied volatility surface in Figure 1 on February 17, 2010.

One-month implied volatilities (Panel A) and the corresponding one-month state prices (Panel B) are shown in Figure 3. Observed bid- and ask implied volatilities are marked with black dots and gray squares.

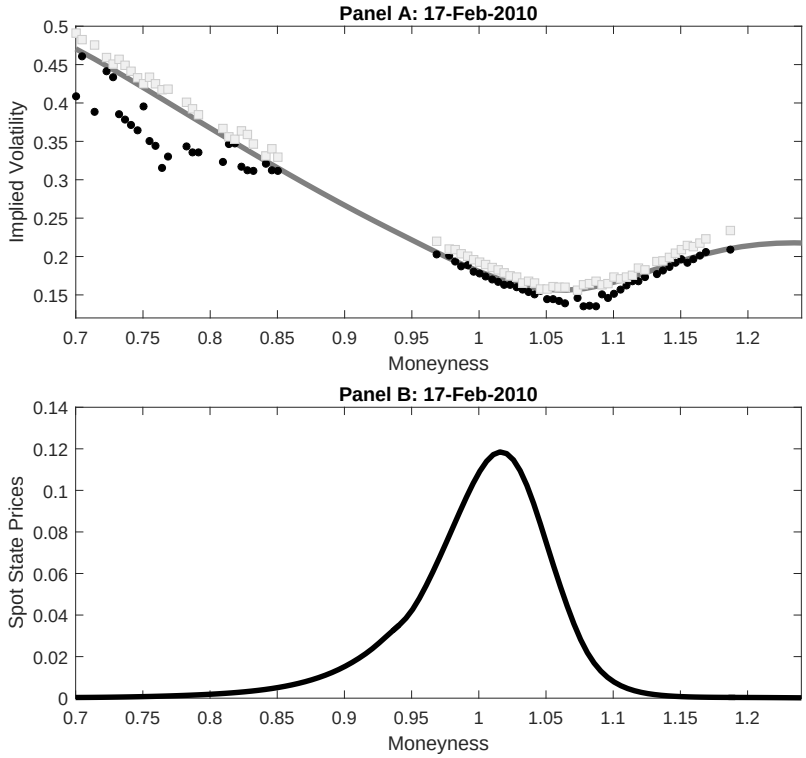

Figure 3. Implied volatilities obtained by our smoothing method (Panel A) and the related spot state prices (Panel B) for a given maturity of 30 days on February 17, 2010. Bid- and ask observed implied volatilities are marked as black dots and gray squares.

B. Transition State Prices

We now explain how we use a given spot state price surface to back out one-month time-homogeneous transition state prices. We utilize the property that multiplying spot state prices $\pi_{0,i}^{(t)}$ for the maturity indexed by t with the one-month transition probability $\pi_{i,j}$ accounts for one possible path of getting to state j in time t^* with corresponding state price $\pi_{0,j}^{(t^*)}$, where t^* indexes the maturity that is 1 month ahead of the maturity indexed by t. Applying the law of total probabilities and

summing over all possible paths i leads to the relation:

$$\pi_{0,j}^{(t^*)} = \sum_{i \in I} \pi_{0,i}^{(t)} \cdot \pi_{i,j} \quad \forall j, t \tag{8}$$

We impose the restriction that $\pi_{i,j} > 0$ and use this system of equations to back out transition state prices $\pi_{i,j}$ in a least squares approach.

An identification of transition state prices by using (8) requires the number of used states N has to be smaller than the number of transitions we apply. We estimate monthly transition state prices out of option prices with a maturity of up to one year. Thus, we are limited to eleven non-overlapping monthly transitions which allows us to only use eleven states to explain the pattern of the spot state price surface and the transition state prices. This demonstrates a problematic feature of implementing the recovery theorem: We need to interpolate the rich implied volatility surface on a rough grid with thirteen points for moneyness (which implies eleven moneyness points for spot state prices) and twelve points for maturity. In addition, we have to make sure that the spot state prices we obtain from this interpolated volatility surface embody tail coverage for high maturity spot state prices as well as fineness for low maturity spot state prices. Ross (2015) already mentions this trade-off in an application of his theorem. One general problem here is that recovered physical probabilities could fail to predict future returns just because of this rough discretization. Also, they do not result in accurate approximations of the full spot state price surface.

Analogously to Audrino et al. (2015), we therefore apply an overlapping approach as our main approach to determine transition state prices.[7]

We recall that the maturity of the full spot state price surface was discretized in steps of 3.5 days with a minimum of 6.5 days and a maximum of 365 days. For each of the maturities at t within 6.5 days and 335.5 days, we link the corresponding spot state prices $\pi_{0,i}^{(t)}$ and spot state prices with a maturity at t plus an additional month $\pi_{0,j}^{(t^*)}$ through one-month transition state prices $\pi_{i,j}$. This results in a total number of 94 (overlapping) transitions and thus a maximum number of

[7] In the main approach, we use the high-dimensional state space to obtain more detailed results. However, we show in the robustness section 6.A that using coarser state spaces do not alter our main results.

N=94 allowed states. We then interpolate the implied volatility surfaces equidistantly so that the corresponding spot state prices are defined on 94 points for the moneyness and again make sure that the discretization hits the current state with moneyness=1.[8]

We back out transition state prices, construct the transition state price matrix Π, and obtain the matrix P= $[p_{i,j}]_{i,j\in I}$ of physical transition probabilities by applying the recovery theorem. Figure 4 illustrates the obtained transition state prices and Figure 5 the corresponding physical transition probabilities. We label this version of recovery **Ross Basic**. The mass of the transition matrix

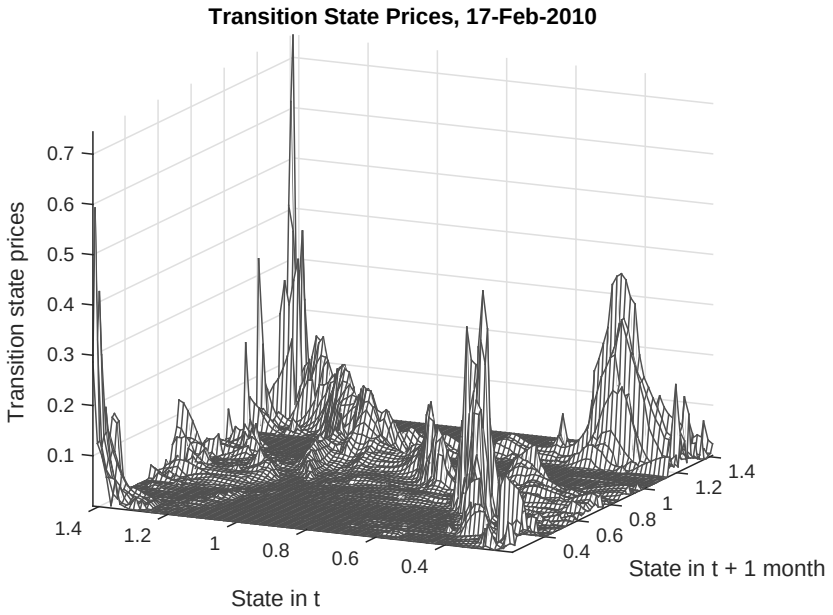

Figure 4. Transition state prices on the 17-Feb-2010 obtained by an implementation without further constraints, only positivity of the transition state prices (basic case).

clusters around the main diagonal. Intuitively, this is not surprising, as the probability of moving from state i to the same state i should be higher than the probability of moving from state i to some far away state j. However, the optimization quite often finds large state prices for states that are far away from the current state. This can happen because option prices (in particular, for the near maturities) are hardly affected by such transition state prices in almost irrelevant states. We try to overcome this problem by including further restrictions that are in line with economic theory.

[8]Note that we labeled the number of states for the fine discretization with N. For ease of notation, we label the number of states that we use for the interpolation (which is smaller than the number of states for the fine discretization) with N as well.

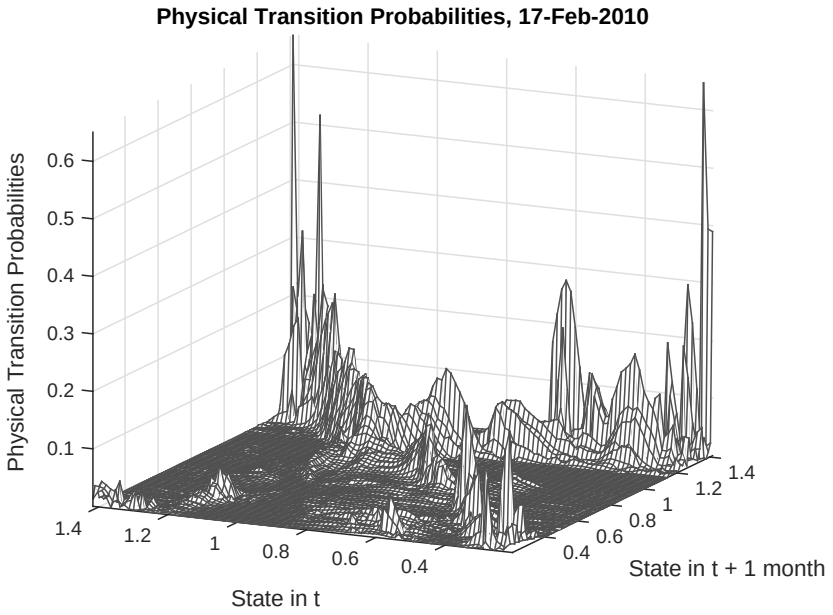

Figure 5. Recovered physical transition probabilities after an application of the recovery theorem (basic case) on the 17-Feb-2010.

We thus apply an additional approach and force the rows in matrix Π to be unimodal with extreme values on the main diagonal. We further demand all rowsums of Π to lie in the interval $[0.9, 1]$. This ensures reasonable values for the state dependent discount factors and thus the state dependent risk-free rates. Figure 6 shows the transition state prices and Figure 7 shows the recovered physical transition probabilities of this modified approach.

By construction, the mass concentrates around the main diagonal without any humps at states that are far away from the initial state. However, for lower (higher) initial states, transition state prices of moving to higher (lower) states are relatively higher than for initial states around the current state. We again think that this is driven by the optimization and the fact that further away states have less importance in explaining option prices. We label this version of recovery **Ross Unimodal**.

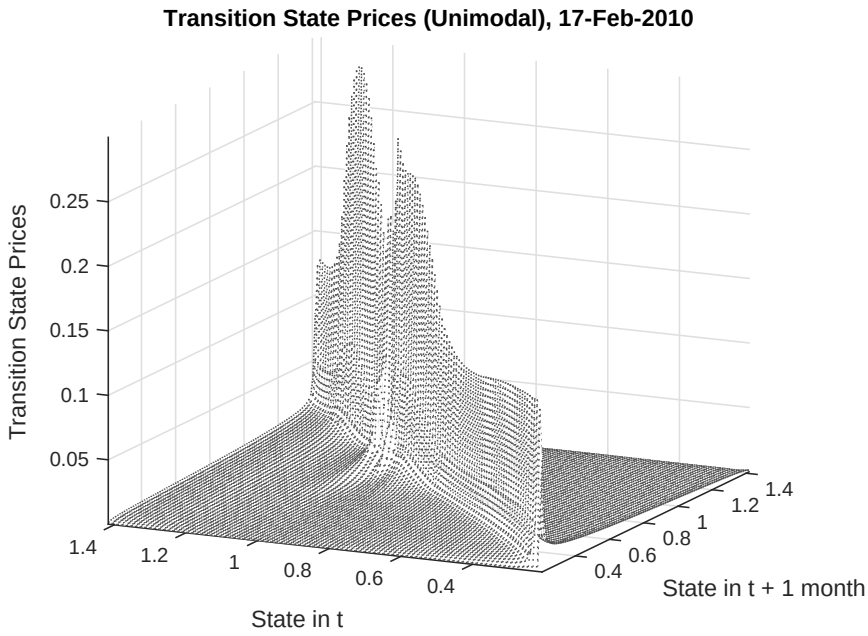

Figure 6. Transition state prices obtained by an implementation with further constraints of unimodality and rowsum bounds (unimodal case) on the 17-Feb-2010.

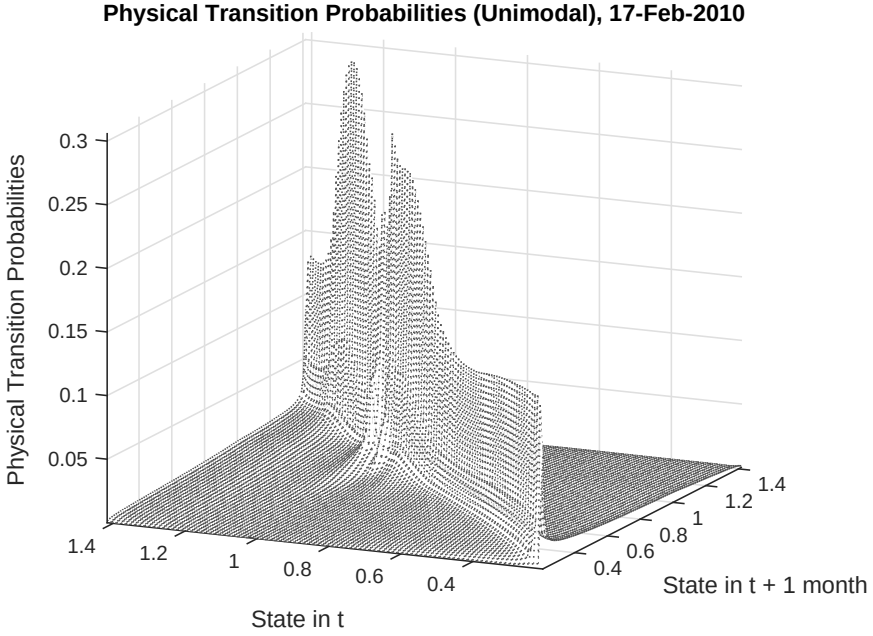

Figure 7. Recovered physical transition probabilities after an application of the recovery theorem (unimodal case) on the 17-Feb-2010.

C. *Recovery without Using Transition State Prices*

The computation of transition state prices appears to be a key challenge in the application of the recovery theorem. Yet there is one row, the one for the current state, where the transition state prices coincide with the spot state prices for one monthly period, which are much easier to estimate. We use this insight to suggest an alternative method.[9] The trick is using the eigenvalue problem in the recovery theorem as in (6) and multiplying both sides from the left with the transition state price matrix Π:

$$\Pi \cdot \Pi z = \Pi \cdot \delta z = \delta(\Pi z) = \delta^2 z \tag{9}$$

Again, the row of Π^2 associated with the current state contains the spot state prices, but now for two months. Additionally, the discount factor δ now appears in the second power to account for the two months. Iterating, we obtain the following relation:

$$\Pi^t z = \delta^t z \tag{10}$$

where t determines how often the transition is being applied. A natural property of the transition state price matrix is that Π^t again represents a transition state prices matrix, but with a transition period that is t times longer than the transition period of Π. The transition state prices of matrix Π^t can again not be inferred directly from option prices. One exception is the row in this matrix that contains transition state prices of moving from the *current* state of the S&P 500 index to any any other state which simply reflect spot state prices $\pi_{0,i}^{(t)}$ with a maturity of t transition periods. We can apply this row property to (10) for different numbers of transition periods t and obtain one equation respectively, which includes known spot state prices with a maturity of t transition periods. We then use this characteristic to identify the pricing kernel implied by z and the utility discount factor δ. In this equation system, the number of unknowns equals to the number of states. A full identification thus requires to use at least as many equations for different maturities t as

[9] See the independent derivation in Jensen et al. (2016).

there are number of states N, which results in the following equation system:

$$\begin{pmatrix} \pi^{(1)}_{0,-N_{low}} & \pi^{(1)}_{0,-N_{low}+1} & \cdots & \pi^{(1)}_{0,N_{high}-1} & \pi^{(1)}_{0,N_{high}} \\ \pi^{(2)}_{0,-N_{low}} & \pi^{(2)}_{0,-N_{low}+1} & \cdots & \pi^{(2)}_{0,N_{high}-1} & \pi^{(2)}_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{(T-1)}_{0,-N_{low}} & \pi^{(T-1)}_{0,-N_{low}+1} & \cdots & \pi^{(T-1)}_{0,N_{high}-1} & \pi^{(T-1)}_{0,N_{high}} \\ \pi^{(T)}_{0,-N_{low}} & \pi^{(T)}_{0,-N_{low}+1} & \cdots & \pi^{(T)}_{0,N_{high}-1} & \pi^{(T)}_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} \frac{z_{-N_{low}}}{z_0} \\ \vdots \\ \frac{z_{-1}}{z_0} \\ 1 \\ \frac{z_1}{z_0} \\ \vdots \\ \frac{z_{N_{high}}}{z_0} \end{pmatrix} = \begin{pmatrix} \delta \\ \delta^2 \\ \vdots \\ \delta^{T-1} \\ \delta^T \end{pmatrix} \tag{11}$$

Appendix C provides more details on how to derive this system of equations. We then solve for δ and $\frac{z_i}{z_0}$ by minimizing the sum of squared errors of these N equations under the constraint that the resulting pricing kernel and the utility discount factor δ have to be non-negative.

To directly recover one-month pricing kernels and one-month physical probabilities $p_{0,j}$, we have to restrict the dimension for moneyness to twelve, as we only can account for a maximum number of twelve transitions and thus twelve maturities. However, we again face the same problem of a rough moneyness grid. [10]

We thus implement an alternative approach that allows for a higher state space dimension. To do so, 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:[11]

$$m^{(t)}_{0,j} = \frac{1}{\delta^{t-1}} \cdot m_{0,j} \tag{12}$$

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. For our application, we choose a transition period of 3.5 days. This allows us to use a fine moneyness grid with 102 points, as the number of maturities within our maturity window of 6.5 to 365 days increases to N=102. We use (12) to

[10] Results and implementation details for this low dimensional approach are given in the robustness section.
[11] Appendix C explains in more detail why this property holds.

find the 1-month pricing kernel from the 3.5 day pricing kernel. Then we can divide the one-month pricing kernel by the one-month state prices to get the one-month physical spot probabilities. We label this approach **Ross Stable**.

D. Testing the Recovered Probabilities.

We test the hypothesis that:

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

For each day t in our sample, we obtain one realization R_t of the monthly S&P 500 future return which is associated with the true pdf p_t. Additionally, we recover the physical spot pdf $\hat{p}_t$ and the corresponding cdf $\hat{P}_t$ for that day. Under the assumption $\hat{p}_t = p_t$, the following transformation of R_t into u_t is i.i.d uniformly distributed, $u_t \sim i.i.d.\ U(0,1)$:[12]

$$u_t = \hat{P}_t(R_t) = \int_{-\infty}^{R_t} \hat{p}_t(x)dx \tag{13}$$

We use four different tests. The Berkowitz (2001) test jointly tests uniformity and the i.i.d. property of u_t. For this test, the series u_t is transformed by applying the inverse standard normal cdf Φ to u_t:

$$z_t = \Phi^{-1}(u_t) = \Phi^{-1}\left(\int_{-\infty}^{R_t} \hat{p}_t(x)dx\right) \tag{14}$$

Under the hypothesis $\hat{p}_t = p_t$, the $z_t \sim i.i.d.\ N(0,1)$. In the Berkowitz setting, independence and standard normality of z_t are tested by estimating the following AR(1) model:

$$z_t - \mu = \rho(z_{t-1} - \mu) + \epsilon_t \tag{15}$$

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

where the null requires $\mu = 0$, Var(ϵ_t)=1 and $\rho = 0$. Berkowitz suggests a likelihood ratio test with

$$LR_3 = -2(LL(0,1,0) - LL(\hat{\mu}, \hat{\sigma}, \hat{\rho})) \tag{16}$$

where LL characterizes the log likelihood of equation (15).

We also apply a test recently introduced by Knüppel (2015). This test first scales the series u_t to $y_t = \sqrt{12}\,(u_t - 0.5)$. In order to test u_t for standard uniformity, the series y_t is tested for scaled uniformity with mean=0 and variance=1. The test then compares the first L moments of the series z_t to the respective moment of its theoretical counterpart in the following GMM-type procedure with test statistic α_L:

$$\alpha_L = T \cdot D_L^{\intercal} \Omega_L^{-1} D_L \tag{17}$$

where D_L is a vector that consists of the differences between the sample moments $\frac{1}{T}\sum_{t=1}^{T} y_t^l$ and the theoretical moments m_l for $l = 1, ..., L$. Ω_L is a consistent covariance matrix estimator of all L respective moment differences between y_t and its theoretical counterpart. We follow Knüppel (2015) and set all elements of Ω_L which represent covariances between odd and even moment differences to zero and apply the test by considering the first four moments (L=4). We account for serial correlation of u_t by estimating a Newey-West covariance matrix with automated lag length as proposed by Andrews (1991).

The Berkowitz test has been used in Bliss and Panigirtzoglou (2004). In that study, the authors argue that it is superior to other possible tests such as the Kolmogorov-Smirnoff (KS) test or the Cramer-van-Mises (CvM) test. While being more powerful than KS and CvM, the Berkowitz test bases its conclusion about standard normality on just the first and the second moment of a series, while ignoring higher moments. The test by Knüppel (2015) has the advantage of testing for higher moments, too, can also deal with autocorrelated data, and still has enough power to test smaller samples. We perform a simulation study with a sample size of $N = 223$ to analyze the power of the Berkowitz, the Knüppel, the Cramer-van-Mises, and the Kolmogorov-Smirnoff test. We find that

the Berkowitz test as well as the test by Knüppel reject the Null hypothesis of standard normality much more likely than the other tests, if we test normally distributed random numbers, but impose a higher or lower variance as in the standard normal case. For our small sample, we therefore rely on the results that we obtain from the Berkowitz test or from the test by Knüppel.

We test several versions of recovery for their ability to predict future returns. As a benchmark, we use the distribution of past five years monthly S&P 500 returns, labeled **Historical Return Distribution**, and a method that assumes a representative investor having a power utility with risk aversion coefficient of 4 for the prediction, labeled **Power Utility with** $\gamma = 4$.[13]

Our sample consists of all dates at which options expire in 30 days, which is on the third Friday of the following month. Starting with the first month, we find the percentile u_t of the future, realized S&P 500 return under the cumulative recovered physical return distribution function.[14] We then shift to the next month and repeat this procedure. This set of points is then tested for uniformity. We apply that testing approach to all recovery versions and to our benchmark models.

IV. Data

We use OptionMetrics to obtain end-of-day option data on the S&P 500 index. We also obtain end-of-day S&P 500 index levels from Datastream. We compute S&P 500 returns from January 1991 to August 2014. Our sample period, as well as our option data, ranges from January 1996 to August 2014. 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 a monthly horizon and end up with N=223 recovered distributions and realized returns.

[13]The results in Bliss and Panigirtzoglou (2004) suggest that a risk aversion factor of 4 in a power utility framework is a reasonable choice.

[14]In our implementation, the values of the recovered cdfs $\hat{P}_t$ are given on a discrete grid, whereas the future realized return does typically not lie on the corresponding grid points. We therefore interpolate the recovered cdf linearly to obtain the transformed points u_t.

V. Empirical Results

We now present the our empirical findings. We show that the Ross recovery theorem is not suitable to make accurate predictions of future returns and does also not imply plausible economic properties. Unless not stated otherwise, all results are discussed for approaches based on a high-dimensional (N=94 if transition state prices are used and $N = 102$ for **Ross Stable**) state space.

A. In-Sample Fit and Predictive Power of Ross' Recovery

We use the p-value of the Berkowitz test as the measure for prediction accuracy. In addition to the prediction accuracy, we are able to analyze the in-sample fit of transition state prices. Therefore, we investigate the corresponding row in the transition state price matrix at which the initial state equals the current state. Transition state prices $\pi_{i,j}$ in this row i should ideally be equal the one-month spot state prices $\pi_{0,j}$. Paradigmatic for **Ross Basic**, Figure 8 shows that this is not the case as the spot state prices (dotted) differ from the corresponding transition state prices (gray line).

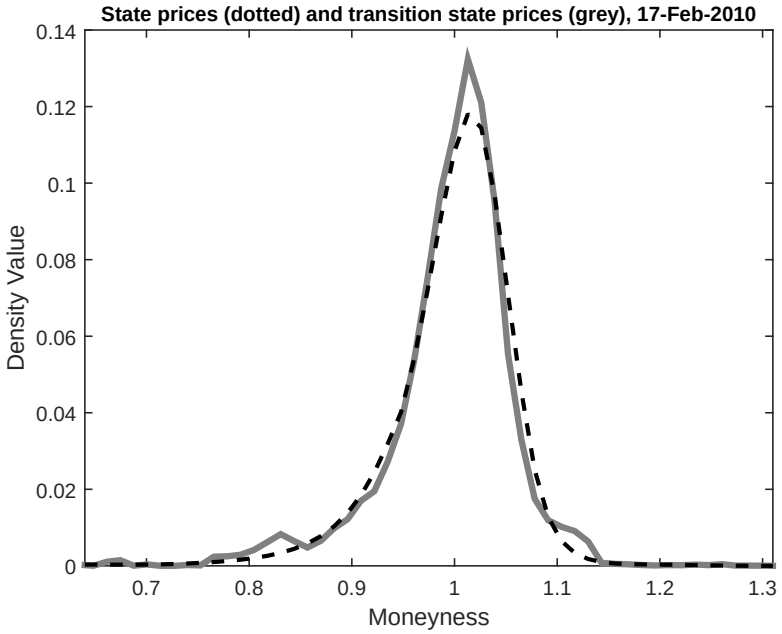

Figure 8. Discrepancy between the state prices (dotted) and the transition state prices with the current state as the initially state (gray line) for **Ross Basic**. Ideally these state price should be the same.

Furthermore, Figure 9 illustrates the different pricing kernels which we obtain if we either divide spot state prices (dotted line) or transition state prices with initial current state (gray line) by recovered physical spot probabilities.

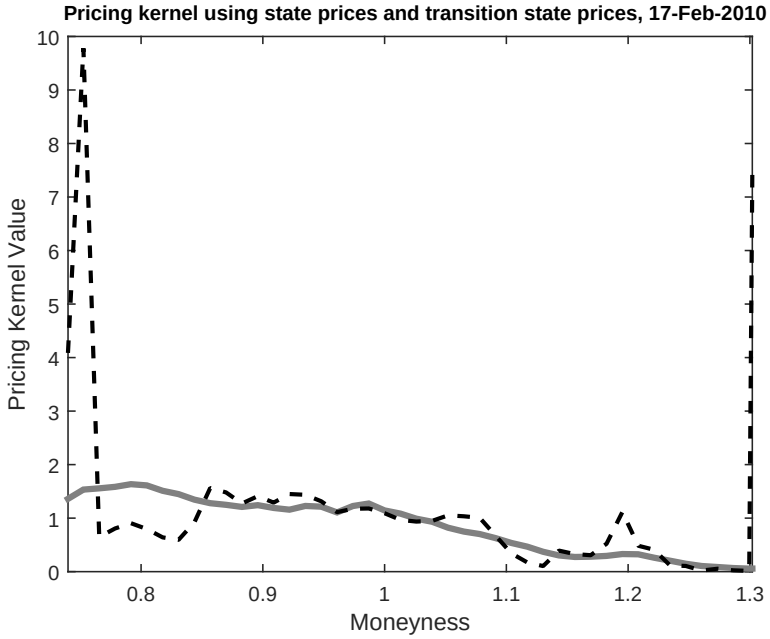

Figure 9. Pricing kernels as the ratio of given spot state prices and recovered physical probabilities (dotted) and as the ratio of transition state prices with initial current state and recovered physical probabilities (gray line) for the basic recovery case.

In that sense, we obtain implied distorted state prices $\hat{\pi}$ which comprise an approximation error for *each* recovery method. For **Ross Stable** and **Power Utility with** $\gamma = 4$ we do not need transition state prices and thus use interpolated spot state prices with state space dimension $N = 102$ for $\hat{\pi}$. Instead of comparing the one-month $\hat{\pi}$ to the true one-month spot state prices π to measure the in-sample fit, we compare the implied volatilities $\hat{\sigma}$ associated with $\hat{\pi}$ to observed implied volatilities associated with π. Appendix D provides details for the computation of $\hat{\sigma}$.

We define the absolute implied volatility error $\epsilon_{l,t}$ as the absolute distance of the approximation $\hat{\sigma}(M_l, t)$ to its corresponding bid- and ask bounds given by $\sigma_{bid}(M_l, t)$ and $\sigma_{ask}(M_l, t)$ for observation l with moneyness M_l and date t:

$$\epsilon_{l,t} = \max\{\sigma_{bid}(M_l, t) - \hat{\sigma}(M_l, t), 0\} + \max\{\hat{\sigma}(M_l, t) - \sigma_{ask}(M_l, t), 0\} \tag{18}$$

As long as $\hat{\sigma}(M_l, t)$ moves within bid- and ask bounds, it is $\epsilon_{l,t} = 0$ and there is no approximation error. We compute the mean of the approximation errors $\epsilon_{l,t}$ over all sample dates and all observations l and obtain the mean absolute implied volatility ($MAIV$) error as a measure for the in-sample performance of each recovery method. Furthermore, we use the p-value of our statistical tests as a measure for forecasting performance.

We test additional implementation methods of the recovery theorem which embody more structure than **Ross Basic**: In **Ross Current State**, we calculate transition state prices as in **Ross Basic** but set transition state prices with initial current state equal to the one-month spot state prices. As we pre-set transition state prices with current initial state for this additional recovery method to spot state prices, we compute their corresponding the $MAIV$ errors based on interpolated spot state prices with a state space dimension of $N = 94$. For **Historical Return Distribution** we do not apply the $MAIV$ as a second performance measure.

Table I shows the p-values and $MAIV$ errors for all introduced methods.

Table I: Mean absolute implied volatility errors and p-values of the Berkowitz test and the test by Knüppel for different recovery methods.

Recovery Method	**MAIV error**	**Berkowitz:** p-value	**Knüppel:** p-value
Ross Basic (I): $\pi_{i,j} > 0$	0.114	0.004	0.017
Ross Current State (II): $\pi_{i,j} > 0$, use spot state prices in Π	0.002	0.003	0.001
Ross Unimodal (III): $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.140	0.045	0.000
Ross Stable (IV): Do not use transition state prices	0.002	0.010	0.010
Power Utility with $\gamma = 4$ (V)	0.002	0.627	0.296
Historical Return Distribution (VI)	-	0.294	0.480

All tests of recovery methods reject our hypothesis at the 5% level with the Berkowitz test as well as the test by Knüppel. **Historical Return Distribution** as well as **Power Utility with** $\gamma = 4$ demonstrate a much better forecasting performance. As expected, the $MAIV$ error is low for methods **(II)** and **(IV)** as it is in that cases based on interpolated spot state prices. For method **(I)** we find this error to be significantly higher. The more structure we impose on transition state prices, the worse we expect the in-sample fit to spot state prices. This explains the even higher $MAIV$ error for method **(III)**.

B. Economic Insights

In theory, the pricing kernel is expected to be positive and monotonically decreasing 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 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.

Figure 4 shows the pricing kernels of different recovery methods for 17-Feb-2010. All methods that are associated with Ross recovery (Panel A-D) deliver non-smooth pricing kernels Among all recovery methods, **Ross Current State** is closest to a monotonically decreasing pricing kernel, while pricing kernels for other recovery methods move rather randomly. **Power Utility with** $\gamma = 4$ implies a monotonically decreasing pricing kernel (Panel E) and is therefore not in line with the pricing kernel puzzle. However, its shape is still more reasonable than the ones we find for almost all of the recovered pricing kernels.

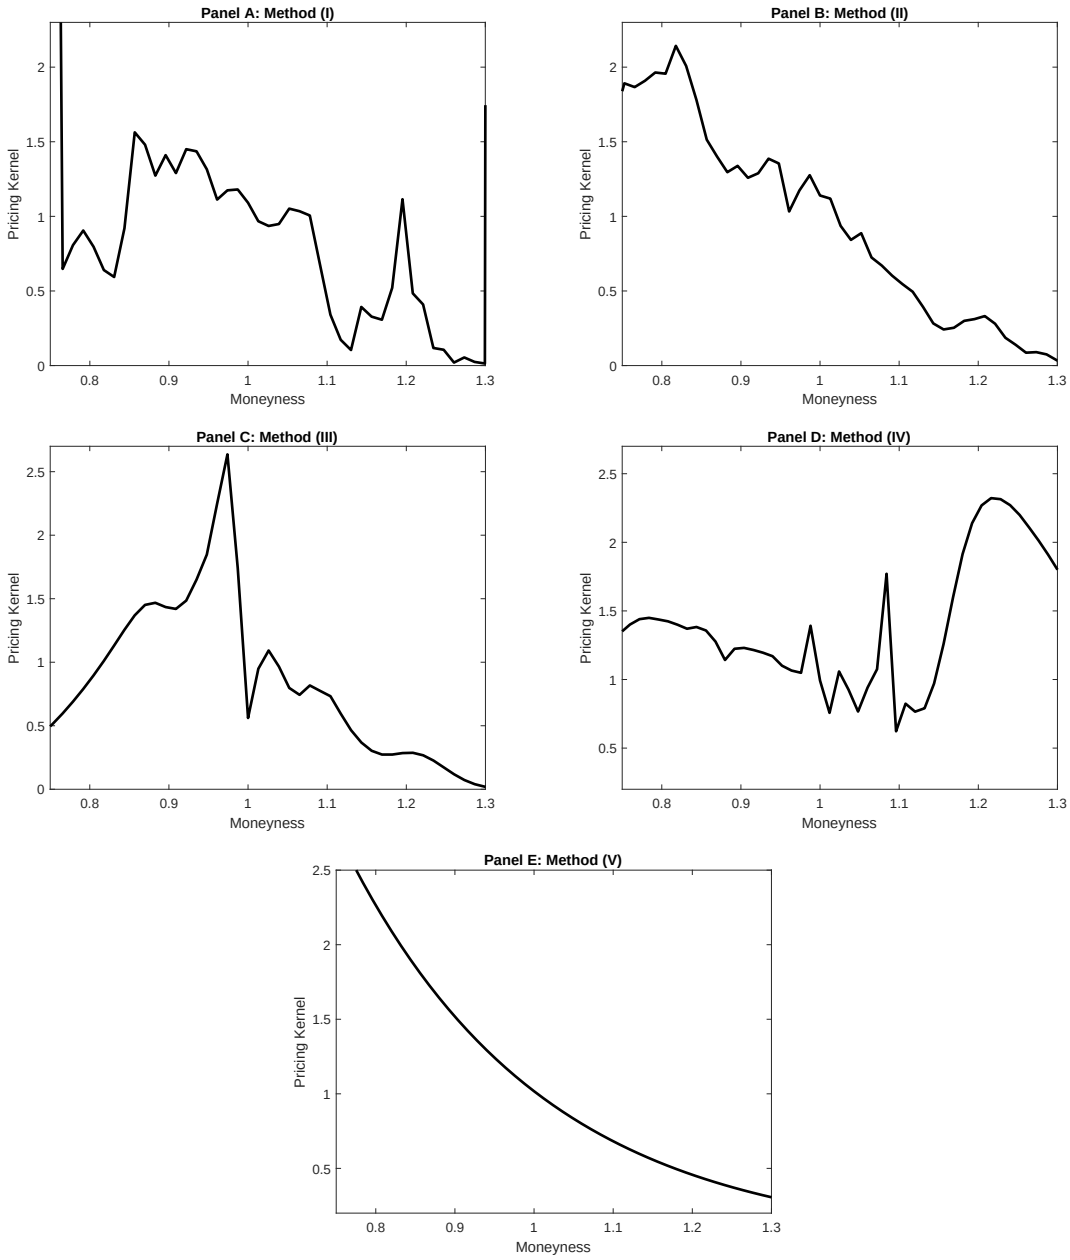

Figure 10. Pricing kernels for different recovery methods on 17-Feb-2010. Panel A shows the pricing kernel for **Ross Basic**, Panel B for **Ross Current State**, Panel C for **Ross Unimodal**, Panel D for **Ross Stable**, and Panel E for **Power Utility with** $\gamma = 4$.

We now document further economic implications of the recovery theorem. We first illustrate why state dependent risk-free rates embody information about the recovered pricing kernel: Each row-sum in the transition state price matrix Π represents a state dependent discount factor. The negative relation of the rowsums of Π and the pricing kernel m denotes that a decreasing pricing kernel, which would be in line with economic theory, would require increasing rowsums for Π. This implies a discount factor (risk free rate) that is positively (negatively) related to the state variable, here the S&P 500 index. We furthermore could determine the pricing kernel by just using information on (state dependent) risk-free rates. Martin and Ross (2013) demonstrate the impact of interest rates on Ross' recovery by showing that the recovered pricing kernel can be inferred from the time series behavior of the long bond and information on the long end of the yield curve.

We further illustrate the significance of interest rates for the recovery theorem from another angle. If we assume that risk-free rates are state-independent, all rowsums of Π would be equal. In Ross' setting, this implies a constant recovered pricing kernel and thus risk-neutrality. We therefore conclude, that the recovery theorem is incompatible with a simple Black-Scholes economy in which risk aversion is possible despite deterministic risk free rates. There is an even more universal problem if general stochastic models are used to explain the dynamics of the state variable. The recovery theorem assumes that option prices contain information on the future return of the underlying. For any given risk-neutral process, which is fully identified by prices of options on the state variable, we can, however, always assume several stochastic processes that describe the behavior of the state variable itself. Thus, the recovery theorem cannot be verified by using synthetic data without any further assumptions.

VI. Robustness

A. Dimension of the State Space

In a simulation study, Tran and Xia (2014) illustrate that the recovered market beliefs are highly dependent on the state space dimension. We now present the implementation and results for the

low dimensional approaches that limit the number of states to $N = 11$ ($N = 12$ for Ross Stable).[15] As we have to deal with the trade-off of tail-coverage for high maturity states and fineness for low-maturity states we implement two different discretization versions for these low-dimensional approaches.

In the first version, we interpolate the volatility surface on thirteen equidistant points for moneyness and twelve monthly points for maturity. We make sure that the whole state space range for moneyness is covered and that the discretization hits the current state at moneyness=1. To allow for a defined state at moneyness=1 is necessary, as we recover physical spot probabilities out of physical transition probabilities with *current* initial state. This first version results in roughly approximated spot state prices for the one-month maturity, but comparably better high-maturity spot state price approximations.

In the second version, we still use an equidistant monthly discretization for the maturity, but apply a non-equidistant discretization for the moneyness with more interpolation points around the at-the-money region. We describe how we set these interpolation points in Appendix B. This version provides a more accurate approximation of the one-month spot state prices, but at the cost of a worse approximation of higher maturity spot state prices.

We apply the same introduced testing procedure to these low-dimensional approaches. We can reject our hypothesis for all recovery methods at the 1% level with p-values close to zero and for both equidistant and non-equidistant low-dimensional discretization versions. As the one-month spot state price approximation is more inaccurate with a small state space dimension, the $MAIV$ errors are also significantly higher than for the high-dimensional approach.

We further want to check if a variation of the state space dimension for our *high-dimensional* approaches changes our test results. So far, we used 94 states for an application of methods (I) - (III). We repeat our study, but reduce the dimensions of the state space by about 10% to 84 states

[15] For models that use transition state prices, we observe, beginning with month one, 11 non-overlapping monthly transitions for one year of data, which limits the state space to $N = 11$. For Ross Stable, we can make use of a maximum of $N = 12$ states instead, as we use spot state prices with 12 different maturities directly.

and by about 20% to 74 states, respectively.[16] Table II provides the results for the original state space, after a 10% reduction of the state space dimension, and after a 20% reduction of the state space dimension.

We observe that changes in the state dimension do empirically not lead to significant changes for the $MAIV$ pricing error as well as for the Berkowitz test results. The historical return distribution is not affected by the discretization scheme and delivers the same p-values independent of the dimension of the state space.

Table II: Mean absolute implied volatility errors and p-values of the Berkowitz test for different recovery methods in the full sample period 01/1996-08/2014 if the dimension of the state space is not reduced (0%), if it is reduced by about 10%, and if it is reduced by about 20%.

Recovery Method	**Reduction (Number) of States**	**MAIV error**	**Berkowitz:** p-value
Ross Basic (I): $\pi_{i,j} > 0$	**0%** (N=94)	0.114	0.004
	10% (N=84)	0.118	0.009
	20% (N=74)	0.114	0.019
Ross Current State (II): $\pi_{i,j} > 0$, use spot state prices in Π	**0%** (N=94)	0.002	0.003
	10% (N=84)	0.002	0.016
	20% (N=74)	0.002	0.017
Ross Unimodal (III): $\pi_{i,j} > 0$ and unimodal, rowsum $\in [0.9, 1]$	**0%** (N=94)	0.140	0.045
	10% (N=84)	0.107	0.022
	20% (N=74)	0.117	0.035
Ross Stable (IV): Do not use transition state prices	**0%** (N=102)	0.002	0.010
	10% (N=92)	0.002	0.015
	20% (N=82)	0.002	0.008
Power Utility with $\gamma = 4$ (V)	**0%** (N=102)	0.002	0.627
	10% (N=92)	0.002	0.649
	20% (N=82)	0.002	0.657

[16] So far, we used 102 states for an application of (IV) and model (V). Here, we reduce the state space dimension by about 10% to 92 states and by about 20% to 82 states.

B. Sub-Samples

Our full sample starts in January 1996 and ends in August 2014 with a sample size of 223 using monthly data. Table III shows the results for the full sample, the sample when excluding the period from 1999 until 2001 related to the Dot-com bubble, and the sample when excluding the period from 2008 until 2009 related to the financial crisis, respectively.

Table III: Mean absolute implied volatility errors and p-values of the Berkowitz test for different recovery methods, if we use data for the full from 01/1996 until 08/2014, if we exclude the period related to the Dot-com bubble, and if we exclude the period related to the financial crisis, respectively.

Recovery Method	**Sample**	**MAIV error**	**Berkowitz:** p-value
Ross Basic (I): $\pi_{i,j} > 0$	Full Sample	0.114	0.004
	No Dotcom bubble	0.116	0.011
	No Financial Crisis	0.114	0.005
Ross Current State (II): $\pi_{i,j} > 0$, use spot state prices in Π	Full Sample	0.002	0.003
	No Dotcom bubble	0.002	0.004
	No Financial Crisis	0.003	0.003
Ross Unimodal (III): $\pi_{i,j} > 0$ and unimodal, rowsum $\in [0.9, 1]$	Full Sample	0.140	0.045
	No Dotcom bubble	0.140	0.072
	No Financial Crisis	0.133	0.036
Ross Stable (IV): Do not use transition state prices	Full Sample	0.002	0.010
	No Dotcom bubble	0.002	0.008
	No Financial Crisis	0.003	0.003
Power Utility with $\gamma = 4$ (V)	Full Sample	0.002	0.627
	No Dotcom bubble	0.002	0.565
	No Financial Crisis	0.003	0.721
Historical Return Distribution (VI)	Full Sample	-	0.294
	No Dotcom bubble	-	0.254
	No Financial Crisis	-	0.766

The statistical tests that we apply are more reliable the larger our sample size is. A large decrease in sample size might therefore itself lead to a change in the test results. We observe a slightly higher p-value for method **(III)** in the absence of the dot-com bubble period and cannot

reject at the 5% level anymore. Even though individual p-values change, the major implications of the test results, however, remain the same.

VII. Conclusion

We implement and test several approaches that recover physical probabilities in the setting of Ross (2015) and find that the corresponding distributions do, according to a Berkowitz test and a test proposed by Knüppel, not equal the empirical distribution of S&P 500 returns. We further present a method that is capable of Ross' recovery without requiring an explicit calculation of transition state prices. The recovered physical probabilities do, however, not match the related empirical distribution as well and cannot be used to predict future S&P 500 returns. On the other hand, simple approaches, such as the empirical distribution of historical five year monthly S&P 500 returns, or a predefined pricing kernel implied by a power utility function with risk aversion parameter $\gamma = 4$, work much better. Further research could focus on applying the recovery theorem to other markets or underlying states. One possible extension is an application to the joint distribution of the S&P 500 returns and the VIX, which is introduced by Jackwerth and Vilkov (2014). Other versions of recovery, that impose less and economically plausible restrictions, are an interesting field for further research as well.

REFERENCES

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

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

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

Bakshi, G., Chabi-Yo, F., and Gao, X. (2015). A Recovery That We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem. *University of Maryland, Ohio State University. Working Paper.*

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. (2015). Misspecified Recovery. *Journal of Finance (forthcoming).*

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state contingent claims implicit in option prices. *The Journal of Business*, 14(3):621–651.

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

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

Cramer, H. (1928). On the Composition of Elementary Errors. *Scandinavian Actuarial Journal.*

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.*

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.*

Jackwerth, J. C. and Vilkov, G. (2014). Asymmetric Volatility Risk: Evidence from Option Market. *University of Konstanz, University of Mannheim. Working Paper, Available at SSRN.*

Jensen, C. S., Lando, D., and Pedersen, Lasse, H. (2016). 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*, 33:270–281.

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

Qin, L. and Linetsky, V. (2015). Positive Eigenfunctions of Markovian Pricing Operators: Hansen-Scheinkman Factorization, Ross Recovery and Long-Term Pricing. *Northwestern University. Working Paper, Available at SSRN.*

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. *The Journal of Finance*, 70(2):615–648.

Schneider, P. and Trojani, F. (2015). (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. Working Paper.*

Walden, J. (2014). Recovery with Unbounded Diffusion Processes. *University of California, Berkeley. Working Paper, Available at SSRN.*

Appendix

Appendix A. Smoothing the Implied Volatility Surface

In order to obtain a smooth volatility surface, we minimize the sum of squared local total second derivatives. We define the total second derivative of the volatility surface $\sigma''_{i,t}$ in the following way:

$$\begin{aligned}\sigma''_{i,t} \;=\; & \frac{\sigma_{i+1,t} - 2\sigma_{i,t} + \sigma_{i-1,t}}{(\Delta i)^2} + \frac{\sigma_{i,t+1} - 2\sigma_{i,t} \;+\; \sigma_{i,t-1}}{(\Delta t)^2} \\ & + \alpha \cdot \frac{\sigma_{i+1,t+1} - \sigma_{i+1,t-1} - \sigma_{i-1,t+1} + \sigma_{i-1,t-1}}{(4\Delta i \Delta t)}\end{aligned} \tag{19}$$

where the first and second terms are the finite differences approximations of the partial second derivative with respect to moneyness and maturity. The third term represents the cross derivative multiplied by a factor α, which controls how much weight is given on this part. We use $\alpha = 1$ in our approach. We evaluate the second derivatives $\sigma''_{i,t}$ at a given maturity on an equidistant grid indexed by t and a given moneyness on an equidistant grid indexed by i, for $i = 1, ..., N$ and $t = 1, ..., T$. Step sizes of moneyness and maturity on the grid are represented by Δi and Δt. We impose the following boundary conditions for all i and t:

$$\begin{gathered}\sigma_{i,0} = \sigma_{i,1}, \quad \sigma_{i,T+1} = \sigma_{i,T} \quad \text{and} \quad \sigma_{0,t} = \sigma_{1,t}, \quad \sigma_{N+1,t} = \sigma_{N,t} \\ \sigma_{0,T+1} = \sigma_{1,T}, \quad \sigma_{N+1,0} = \sigma_{N,1} \quad \text{and} \quad \sigma_{0,0} = \sigma_{1,1}, \quad \sigma_{N+1,T+1} = \sigma_{N,T}\end{gathered} \tag{20}$$

In the fast and stable method by Jackwerth (2004), a trade-off parameter λ balances between smoothness and a small deviation from mid prices. In this approach, we instead allow the implied volatilities to move freely within observed bid- and ask implied volatility bounds and penalize them if they move outside of these bounds. Therefore, we define the penalty terms $\delta_l^{ask} = \max\left(\sigma_{i_l,t_l} - \sigma_{i_l,t_l}^{ask}, 0\right)$ and $\delta_l^{bid} = \max\left(\sigma_{i_l,t_l}^{bid} - \sigma_{i_l,t_l}, 0\right)$. The combined penalty terms $\delta_l^{ask} + \delta_l^{bid}$ signify how far implied volatilities moves outside the bid- and ask bounds. We optimize for smoothness

and weight the sum of all squared combined penalty terms with a trade-off parameter λ:

$$\begin{aligned} \min_{\sigma_{i,t},\ \delta_l^{\text{bid}},\ \delta_l^{\text{ask}}} \quad & \frac{1}{TN} \cdot \sum_{t=1}^{T} \sum_{i=1}^{N} (\sigma''_{i,t})^2 \quad + \quad \lambda \cdot \frac{1}{L} \cdot \sum_{l=1}^{L} \left(\delta_l^{\text{bid}} + \delta_l^{\text{ask}} \right)^2 \\ & s.t. \quad \delta_l^{ask} \ \geq \ \sigma_{i_l,t_l} - \sigma_{i_l,t_l}^{ask} \\ & \qquad \delta_l^{bid} \ \geq \ \sigma_{i_l,t_l}^{bid} - \sigma_{i_l,t_l} \\ & \qquad \sigma_{i,t},\ \delta_l^{\text{bid}},\ \delta_l^{\text{ask}},\ \geq \ 0 \end{aligned} \tag{21}$$

We start with a large λ which means that observed implied volatilities move within or close to bid- or ask bounds. Iteratively, we decrease λ until we obtain smooth and strictly positive state prices. We consider a spot state price surface as smooth, when either the absolute or the relative change of spot state prices from one state i to the next state $i+1$ is not larger (not smaller) than the (negative) change in moneyness if i is larger (smaller) than $i^\star$, where $i^\star$ is associated with the largest spot state price $\pi^\star$. Thus, we push the state prices to be unimodal, which avoids spurious probability masses in the (far) tails.

The decrease of λ leads to a flat volatility surface in an extreme case because we give less weight to the penalty for observed implied volatilities that move outside bid- and ask bounds. As a flat volatility surface satisfies the required conditions of smoothness and positivity for state prices, they will ultimately be fulfilled with our iterative procedure.

Appendix B. Non-Equidistant Interpolation of the Spot State Price Surface

Our goal is to interpolate the original state price surface with twelve monthly discretization points across maturity and eleven discretization points across moneyness. We first fix five interpolation (IP) points for moneyness levels of spot state prices: One point each at the beginning (point **p1**) and the end (point **p11**) of the whole state space range, one point at the beginning (point **p3**) and the end (point **p9**) in which the one-month spot state prices have non-zero values, and one point at the current state with moneyness =1. We optimize for the location of the six remaining points in a way that the interpolated volatility surface fits the original volatility surface best in a least

squares sense, while where one IP point (point **p2**) lies in between **p1** and **p3**, four IP points lie in between **p3** and **p9**, and one IP point (point **p10**) lies in between **p9** and **p11**.

Appendix C. Recovery without Using Transition State Prices

We start with equation (10) and show how it can be used to achieve recovery without using transition state prices in our empirical approach. Recall that the $0-th$ row in Π represents the spot state prices with a maturity equal to the time period of the transition. The $0-th$ row of Π^t then represents the spot state prices with a maturity equal to the time period of t transitions. We can apply this information to equation (10). We first express this equation system as follows:

$$\begin{pmatrix} \pi^{(t)}_{-N_{low},-N_{low}} & \pi^{(t)}_{-N_{low},2} & \cdots & \pi^{(t)}_{-N_{low},N_{high}-1} & \pi^{(t)}_{-N_{low},N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{(t)}_{0,-N_{low}} & \pi^{(t)}_{0,-N_{low}+1} & \cdots & \pi^{(t)}_{0,N_{high}-1} & \pi^{(t)}_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{(t)}_{N_{high},1} & \pi^{(t)}_{N_{high},2} & \cdots & \pi^{(t)}_{N_{high},N_{high}-1} & \pi^{(t)}_{N_{high},N_{high}} \end{pmatrix} \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix} = \delta^t \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix} \tag{22}$$

where $\pi^{(t)}_{i,j}$ is the transition state price of moving from state i to state j in a time equal to t transition periods. The corresponding $0-th$ equation does only contain information about the spot state prices, which can be inferred directly from option data:

$$\begin{pmatrix} \pi^{(t)}_{0,-N_{low}} & \pi^{(t)}_{0,-N_{low}+1} & \cdots & \pi^{(t)}_{0,N_{high}-1} & \pi^{(t)}_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} z_{-N_{low}} \\ \vdots \\ z_0 \\ \vdots \\ z_{N_{high}} \end{pmatrix} = \delta^t \cdot z_0 \tag{23}$$

After dividing both sides by z_0, we apply this equation for every $t = 1, ..., T$ and obtain the following equation system:

$$\begin{pmatrix} \pi^{(1)}_{0,-N_{low}} & \pi^{(1)}_{0,-N_{low}+1} & \cdots & \pi^{(1)}_{0,N_{high}-1} & \pi^{(1)}_{0,N_{high}} \\ \pi^{(2)}_{0,-N_{low}} & \pi^{(2)}_{0,-N_{low}+1} & \cdots & \pi^{(2)}_{0,N_{high}-1} & \pi^{(2)}_{0,N_{high}} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \pi^{(T-1)}_{0,-N_{low}} & \pi^{(T-1)}_{0,-N_{low}+1} & \cdots & \pi^{(T-1)}_{0,N_{high}-1} & \pi^{(T-1)}_{0,N_{high}} \\ \pi^{(T)}_{0,-N_{low}} & \pi^{(T)}_{0,-N_{low}+1} & \cdots & \pi^{(T)}_{0,N_{high}-1} & \pi^{(T)}_{0,N_{high}} \end{pmatrix} \cdot \begin{pmatrix} \frac{z_{-N_{low}}}{z_0} \\ \vdots \\ \frac{z_{-1}}{z_0} \\ 1 \\ \frac{z_1}{z_0} \\ \vdots \\ \frac{z_{N_{high}}}{z_0} \end{pmatrix} = \begin{pmatrix} \delta \\ \delta^2 \\ \vdots \\ \delta^{T-1} \\ \delta^T \end{pmatrix} \tag{24}$$

With $\frac{z_j}{z_0}$, $j \in I$ and δ, we obtain N unknowns in this nonlinear equation set. As δ is the utility discount factor, we limit it with an upper bound of 1 and solve the equation system by means of least squares. We force the ratios $\frac{z_j}{z_0}$ to be non-negative, as this ratio is directly linked to the pricing kernel. We then directly compute the pricing kernel by using its relation to the eigenvector z. Equation (10) implies that the pricing kernel can be defined for different maturities of t transition periods in a way that its structure will be time independent and it will only differ in the discount factor:

$$m^{(t)}_{0,j} = \frac{1}{\delta^t} \cdot \frac{z_0}{z_j}, \quad j \in I \tag{25}$$

where $m^{(t)}_{0,j}$ represents the pricing kernel value for state j with a maturity of t transition periods. We can use this property and link pricing kernels with different maturities in the following way:

$$m^{(t)}_{0,j} \quad = \quad \frac{1}{\delta^t} \cdot \frac{z_0}{z_j} \quad = \quad \frac{1}{\delta^{t-1}} \cdot \frac{1}{\delta} \cdot \frac{z_0}{z_j} \quad = \quad \frac{1}{\delta^{t-1}} \cdot m_{0,j}, \quad j \in I \tag{26}$$

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

Appendix D. In-Sample Pricing Error for Recovery

In order to compare the pricing error that results from a corresponding recovery method, we use information on spot state prices that are implied by the recovery method $\hat{\pi}$ and true spot state prices π. We therefore compute one-month option prices $\hat{C}$ that stem from the distorted state prices $\hat{\pi}$ by numerical integration:

$$\hat{C}(M_l) = \int_0^\infty \hat{\pi}_t(S) \cdot \max(S - M_l, 0) dS \tag{27}$$

for every moneyness level M_l at which we observe an one-month option price. We then transform option prices $\hat{C}(M_l)$ into implied volatilities $\hat{\sigma}(M_l)$. We do this for each sample date t and specify implied volatilities as $\hat{\sigma}(M_l, t)$.