



Universität St.Gallen

An Empirical Analysis of the Ross Recovery Theorem

Francesco Audrino, Robert Huitema, Markus Ludwig

May 2014 Discussion Paper no. 2014-11

Editor: Martina Flockerzi
University of St.Gallen
School of Economics and Political Science
Department of Economics
Bodanstrasse 8
CH-9000 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 31 35
Email seps@unisg.ch

Publisher: School of Economics and Political Science
Department of Economics
University of St.Gallen
Bodanstrasse 8
CH-9000 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 31 35

Electronic Publication: <http://www.seps.unisg.ch>

An Empirical Analysis of the Ross Recovery Theorem¹

Francesco Audrino, Robert Huitema, Markus Ludwig

Author's address:

Prof. Dr. Francesco Audrino
University of St. Gallen
Bodanstrasse 6
CH-9000 St. Gallen
Phone +41 71 224 24 31
Fax +41 71 224 28 94
Email francesco.audrino@unisg.ch

Robert Huitema
University of Zurich
Plattenstrasse 14
CH-8032 Zurich
Email robert.huitema@bf.uzh.ch

Markus Ludwig
University of Zurich
Plattenstrasse 14
CH-8032 Zurich
Email markus.ludwig@bf.uzh.ch

¹ We acknowledge discussions with Peter Carr, Andrew Jeffrey, Markus Leippold, Eckhard Platen, Stephen Ross and Liuren Wu. Part of this work was done while Markus Ludwig was visiting the Finance Discipline Group at the University of Technology, Sydney. Corresponding author: Markus Ludwig. E-mail: markus.ludwig@bf.uzh.ch

Abstract

Building on the results of Ludwig (2012), we propose a method to construct robust time-homogeneous Markov chains that capture the risk-neutral transition of state prices from current snapshots of option prices on the S&P 500 index. Using the recovery theorem of Ross (2013), we then derive the market's forecast of the real-world return density and investigate the predictive information content of its moments. We find that changes in the recovered moments can be used to time the index, yielding strategies that not only outperform the market, but are also significantly less volatile.

Keywords

Risk-neutral density, real-world density, pricing kernel, risk aversion, predictive information.

JEL Classification

C14, C58, G13.

Some moments are nice, some are nicer, some are even worth writing about.

– Charles Bukowski

Recently, Ross (2013) has shown that the market's risk aversion, in the form of a transition independent pricing kernel, can be recovered from the risk-neutral transition matrix of a Markovian state variable. Knowledge of the pricing kernel allows us to obtain the market's subjective assessment of real-world probabilities from risk-neutral densities, which makes the information embedded in option prices directly accessible to applications such as risk management, portfolio optimization and the design of trading strategies. Ross' recovery theorem is intriguing because, in contrast to previous literature, it does neither rely on historical returns nor restrictions on the shape of the pricing kernel. His work also challenges conventional wisdom. While implied volatility has long been used to gauge the market's perception of risk, option prices were considered to be silent when it comes to predicting the actual return, let alone the entire real-world distribution.

The recovery theorem has already prompted several theoretical extensions. Carr and Yu (2012) investigate the assumptions behind Ross' approach and propose an alternative for diffusions on a bounded state space, with restrictions placed on the form and dynamics of the numeraire portfolio instead of the preferences of an agent. Dubynskiy and Goldstein (2013) investigate implications of bounds on the state space and theorize that their specific placement might give rise to econometric fragility. Martin and Ross (2013) extend the results to fixed income markets and elucidate the long-term implications of Ross' assumptions. Huang and Shaliastovich (2013) consider an extension that accounts for a preference for the early resolution of uncertainty, but fall back on historical data.

We set out to investigate whether or not the recovery theorem yields predictive information beyond what can be gleaned from the moments of risk-neutral densities. Our work contributes to the existing literature along several dimensions. We present the first comprehensive empirical analysis of the recovery theorem, using options on the S&P 500 spanning a period of 13 years. Building on the results of Ludwig (2012), we present an algorithm to construct a time-homogeneous Markov chain from current option prices. We also address the robustness of our results and investigate what drives the pricing kernel.

The remainder of this paper is organized as follows. The next section outlines how the relation between option prices and the real-world has evolved and then briefly introduces the key ideas behind Ross recovery. Section 2 provides a more detailed look at the individual steps involved in applying the recovery theorem in practice. In Section 3 we present the results of our empirical study. Section 4 sketches out potential extensions and concludes.

1 Option Prices and the Real World

Literature on option pricing has traditionally focused on specifying the stochastic process of an asset and deriving option prices following from the imposed dynamics and no-arbitrage. By equating the price of a contingent claim to that of a dynamic trading strategy, involving only the underlying and a cash account, pricing is possible without knowledge of the expected return on the underlying. The framework of risk-neutral valuation was established by Black and Scholes (1973) and Merton (1973). Ever since,

numerous extensions have sought to improve on the specification of the stochastic process, so as to better match the market prices of options. Prominent examples include stochastic volatility models, e.g., Hull and White (1987), Stein and Stein (1991) and Heston (1993), models with jump processes, e.g., Merton (1976), Bates (1991) and Carr et al. (2002), and combinations of stochastic volatility and jumps, e.g., Bates (1996) and Bates (2000).

With the increasing liquidity of markets for contingent claims, options have become assets in their own right, to the extent that market prices now determine the model parameters through the inverse problem of calibration. The quality of a model is often assessed in the space of implied volatilities, which highlights deviations from the assumptions underlying the Black-Merton-Scholes (BMS) framework and facilitates comparisons across different strikes and maturities. Despite frequent re-calibration, traditional models are typically not able to fully capture the rich patterns observable in market prices. This has led to a somewhat orthogonal approach that seeks to provide consistent quotes in a specified set, given the price of the underlying asset and option prices at several strikes and maturities, without specifying the dynamics of the underlying. This literature encompasses work on implied volatility, price, and state price density surfaces. Among the proposed methods are polynomials, e.g., Dumas, Fleming, and Whaley (1998), kernel smoothing, e.g., Aït-Sahalia and Lo (1998), regression trees, e.g., Audrino and Colangelo (2010), B-splines, e.g., Fengler and Hin (2013), and neural networks, e.g., Ludwig (2012).

Over the last couple of years, there has been increased interest in the predictive information content of market option prices. Option implied

volatility has long been used as an important sentiment indicator, the most notable example being the Chicago Board of Options Exchange Volatility Index (VIX), often referred to as the investor fear gauge, cf. Whaley (2000). Implied volatility has also been demonstrated to be an upward biased predictor of future realized volatility, see e.g., Canina and Figlewski (1993). More recently, research has focused on the predictive information content of the implied volatility skew, e.g., Xing, Zhang, and Zhao (2010) and Ratcliff (2013). Higher order moments of the risk-neutral distribution, based on the results in Bakshi, Kapadia, and Madan (2003), have been investigated in Rehman and Vilkov (2012) and Conrad, Dittmar, and Ghysels (2013). Bollerslev, Tauchen, and Zhou (2009), Goyal and Saretto (2009) and Drechsler and Yaron (2011) find that differences between risk-neutral and realized measures of volatility predict future returns.

The mapping from risk-neutral to actual probabilities has also been studied in the literature on the stochastic discount factor or pricing kernel, defined as the ratio between the two measures, projected onto return states. Prior to the recovery theorem, the real-world density had to be estimated by resorting to time-series data, see e.g., Jackwerth (2000), Ait-Sahalia and Lo (2000), Rosenberg and Engle (2002) and Bliss and Panigirtzoglou (2004). Such estimates are not only highly sensitive to the chosen time period, but by nature of construction less expressive than an estimate based on current quotes. The mismatch between a density capturing the rich amalgam of current investor sentiment and one based on decades of historical data invariably manifests in the pricing kernel. Features deviating from the monotonically decreasing shape prescribed by neoclassical theory were perceived as a puzzle and have sparked a rich literature, see e.g., Ziegler (2007) and Hens and Reichlin (2013).

The Recovery Theorem

Ross (1976) and Cox and Ross (1976) demonstrate that in a dynamically complete arbitrage-free market, the price of an option is given by the expected present value of the payoff under the risk-neutral density (RND). The RND is the continuous-state counterpart to the prices of Arrow-Debreu state-contingent claims, that pay \$1 in *one* ending state and nothing in all other states. Extending these results, Banz and Miller (1978) and Breeden and Litzenberger (1978) show that if we know the prices of European options for all possible levels of the underlying, the discounted RND at expiration is equal to the second derivative of the option prices with respect to strike. Two additional conditions allow Ross (2013) to uniquely disentangle risk aversion from what the market believes the real-world probabilities to be: (i) the risk-neutral process evolves as a discrete time-homogeneous Markov chain on a finite state space, (ii) the pricing kernel is path independent, and therefore, utility is not state dependent.

Consider a time-homogeneous process X_t, X_{t+1}, \dots on a finite state space with values in $1, \dots, n$. Since calendar time is irrelevant, the transition probability of moving from state i at time t to state j at time $t + 1$ is given by:

$$P_{i,j} = Pr(X_{t+1} = j | X_t = i) \tag{1}$$

where $P_{n \times n}$ denotes the one-step ahead transition matrix and is clearly element-wise non-negative. We also assume that P is irreducible such that $P^t > 0$ for some t , that is, all states can be reached from all other states. From the Perron-Frobenius theorem we know that a non-negative irreducible matrix has a unique positive eigenvalue λ and corresponding dominant left-

and right-eigenvectors that are both unique and strictly positive, cf. Meyer (2000).

$$\mathbf{v}^\top P = \lambda \mathbf{v}^\top \quad (2)$$

$$P \mathbf{z} = \lambda \mathbf{z} \quad (3)$$

Here \mathbf{v} corresponds to the long-run distribution of the Markov chain up to scale, and \mathbf{z} traces the variations in scale over starting states. If the rows of P sum to one, it is a stochastic matrix, and \mathbf{z} is a constant. For the recovery theorem, P is substochastic, as it captures the dynamics of the discounted RND, that is, the state prices.

The elements $P_{i,j}$ of the state price transition matrix correspond to the prices of single period Arrow-Debreu securities, indexed by the starting and ending state. The price of an Arrow-Debreu security can be decomposed as:

$$P_{i,j} = \delta M_{i,j} F_{i,j} \quad (4)$$

where δ is the market's average discount rate, $M_{n \times n}$ is the pricing kernel, and $F_{n \times n}$ denotes the real-world transition matrix. It is apparent that knowledge of P is not sufficient to uniquely determine the right-hand side.

However, the assumption that the pricing kernel is transition independent means that it only depends on the marginal rate of substitution between the future and current consumption. Defining the vector of marginal utilities $\mathbf{d} = (d_1, \dots, d_n)^\top$ allows us to write (4) as:

$$P_{i,j} = \delta \frac{d_j}{d_i} F_{i,j} \quad (5)$$

Defining the diagonal matrix $D_{n \times n}$ with \mathbf{d} on the main diagonal allows us to rearrange for F as follows:

$$\begin{aligned} P &= \delta D^{-1} F D \\ F &= \delta^{-1} D P D^{-1} \end{aligned} \tag{6}$$

Since F is stochastic, we have $F\mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ is a vector of ones. Combined with (6) we obtain:

$$P D^{-1} \mathbf{1} = \delta D^{-1} \mathbf{1} \tag{7}$$

If we define the vector \mathbf{z} to contain the inverse of the diagonal elements of D , we obtain the unique solution from (3), with λ corresponding to the discount factor δ .

To summarize, time-homogeneity and irreducibility make sure that the eigenvalue problem has only one positive solution, and since \mathbf{z} corresponds to the inverse of the pricing kernel, which by no-arbitrage constraints must be strictly positive, Ross (2013) is able to achieve recovery.

2 Steps to Recovery

The results of Breeden and Litzenberger (1978) only allow us to compute the prices for Arrow-Debreu claims contingent on the current state of the underlying. The construction of a full transition matrix, including claims conditional on starting states that differ from the current one, is thus at the core of applying the recovery theorem.

From Option Prices to State Prices

In reality we don't observe a continuum of traded strikes, so we first have to estimate option prices on a dense grid. Market prices also contain microstructure noise, which poses a challenge, since derivatives amplify irregularities. Furthermore, in order to estimate the transition matrix, we need a current snapshot of the entire price surface, spanning multiple maturities over a uniform state space. These requirements alone rule out the vast majority of methods proposed to estimate option implied densities, as they either aggregate data over time, deal only with single maturities, or neglect the issue of extrapolation.¹

It is our goal to fully capture the rich information priced into options by means of supply and demand. In order to avoid the risk of model misspecification, we follow the neural network approach of Ludwig (2012), which has been shown to yield robust state price surfaces that are both highly flexible and free of static arbitrage.

Just like linear models, neural networks approximate functions as a linear combination of features. However, instead of fixed nonlinear transformations, the form of the basis expansions is learned at the same time as the coefficients of the simple linear model. Consider modeling implied volatility σ for a fixed maturity τ as a function of strike x , using polynomials of up to second order. The resulting function is additive in the features derived

¹See e.g., Shimko (1993), Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (1998), Bliss and Panigirtzoglou (2002), Ait-Sahalia and Duarte (2003), Figlewski (2010) and Fengler and Hin (2013).

from the basis expansions $g_i(x) = x^i$ and can be described as:

$$\sigma(x) = \sum_{i=0}^2 \beta_i \cdot g_i(x)$$

where β_i corresponds to the coefficient of the i^{th} feature. Now, consider a neural network that models total implied variance $\nu(x, \tau) := \sigma^2(x, \tau) \cdot \tau$ as a function of both strike and time to maturity:

$$\begin{aligned} \nu(x, \tau) &= \sum_{i=1}^H \beta_i \cdot h_i(x, \tau) & (8) \\ h_i(x, \tau) &= \frac{1}{1 + e^{-c_i(x, \tau)}} \\ c_i(x, \tau) &= \alpha_{0i} + \alpha_{1i}x + \alpha_{2i}\sqrt{\tau} \end{aligned}$$

Here, the basis function $h_i(x, \tau)$ is a logistic sigmoid. The weights α_{1i} and α_{2i} control the steepness and thus modulate the amount of nonlinearity introduced by each expansion. The constant α_{0i} has the effect of shifting the sigmoid horizontally, allowing the total variance surface to be build up additively by regions of varying curvature.²

Given a sufficient number of basis expansions H , neural networks can approximate any continuous function on a compact input domain to an arbitrary degree of accuracy, see e.g., Hornik, Stinchcombe, and White (1989).

However, the usefulness of an estimator depends on its ability to successfully generalize beyond observable data. The real challenge is thus to find a model that is flexible enough to capture the relationships implicit in a set of sparse and noisy observations without memorizing them.

²The intercept β_0 is omitted for notational clarity.

In neural networks, model complexity can be controlled both through the number of basis expansions and through regularization, that is, a penalty on the norm of α and β . Since neural networks have to be trained using iterative numerical methods, a comparable effect can be achieved by stopping the optimization before convergence, as this keeps the weights close to their highly regularized initial values, see e.g., Bishop (2006).

A common approach to determine the optimal model complexity is cross-validation, which aims to gauge the generalization performance of an estimator by omitting some observations during training. However, due to the irregular data design of exchange traded equity options, cross-validation is of limited usefulness, as it can only check locations for which we actually have market quotes.

That means we cannot assess the quality of a surface estimator between maturities, or for the sizable regions outside of traded strikes. The key idea of Ludwig (2012) is to use two non-trivial problems, namely extrapolation and no-arbitrage constraints, to solve a third, namely finding the optimal model complexity of neural networks.

This is done by generating a population of estimators from networks with a varying number of basis expansions and random starting parameters. The neighborhoods of these initial solutions are then explored by running a regularized gradient descent algorithm for only 10 steps. The estimators are then checked for arbitrage violations. Due to the considerable amount of extrapolation involved in obtaining surfaces that extend into the tails of the density, these checks allow us to eliminate solutions that exhibit inferior generalization performance. Within the subset of valid networks, solutions can be ranked based on the empirical loss. This method yields an estimator

that integrates the information contained in multiple options series into a globally consistent and arbitrage-free surface.

From State Prices to Markov Chains

Given a matrix of state prices $S_{n \times m}$ for m equidistantly spaced maturities, Ross (2013) shows that if $m \geq n$, we can estimate P because it specifies the time-homogeneous transition from one maturity to the next:

$$S_{:,t+1}^\top = S_{:,t}^\top P, \quad t = 1, 2, \dots, m-1 \quad (9)$$

Let $A^\top = S_{:,1:m-1}$ and $B^\top = S_{:,2:m}$ contain state prices, equation (9) gives rise to a matrix factorization problem:

$$\min_{P \geq 0} \|AP - B\|_2^2 \quad (10)$$

which is equivalent to n independent vector least squares problems. Let $\mathbf{p}_j := P_{:,j}$ and $\mathbf{b}_j := B_{:,j}$ we have:

$$\min_{\mathbf{p}_j \geq 0} \|\mathbf{A}\mathbf{p}_j - \mathbf{b}_j\|_2^2, \quad j = 1, 2, \dots, n \quad (11)$$

This non-negative least squares problem can be solved with standard algorithms, cf. Lawson and Hanson (1974). However, active-set methods depend on $A^\top A$, which renders them infeasible if A is ill-conditioned, as is the case for our application.³ Considering that the transition matrix is at the center of the recovery theorem, and given that the estimation of state prices from option prices is already ill-posed, a high sensitivity of solutions

³The average condition number for $A_{11 \times 11}$ is $2.17 \cdot 10^8$, for an overdetermined higher resolution system $A_{621 \times 150}$ it is $6.82 \cdot 10^{12}$.

for P with respect to small perturbations in A is highly undesirable. We therefore consider the damped least squares problem:

$$\min_{\mathbf{p}} \quad \|A\mathbf{p} - \mathbf{b}\|_2^2 + \zeta \|\mathbf{p}\|_2^2 \quad (12)$$

where the scalar ζ controls the trade-off between fit and stability of the solution by reducing the impact of small singular values. This technique to stabilize the solution of inverse problems is known as ridge regression or Tikhonov regularization, cf. Tikhonov and Arsenin (1977). The regularized formulation is also a least squares problem. Augmenting the constrained least squares problem yields:

$$P = \operatorname{argmin}_{P \geq 0} \quad \sum_{j=1}^n \left\| \begin{bmatrix} A \\ \zeta I \end{bmatrix} \mathbf{p}_j - \begin{bmatrix} \mathbf{b}_j \\ \mathbf{0} \end{bmatrix} \right\|_2^2 \quad (13)$$

where I denotes an identity matrix and $\mathbf{0}$ is a vector of zeros. A natural approach to determine ζ is to minimize the discrepancy between observable state prices and the *unrolled* Markov chain. Let $\boldsymbol{\nu}_i$ be a vector with 1 in the i^{th} position and zeros elsewhere, the t -step ahead state prices implied by a given transition matrix can be computed as:

$$U_{:,t}^\top = \boldsymbol{\nu}_i^\top P^t, \quad t = 1, 2, \dots, m \quad (14)$$

Since we define our state space symmetrically around the current state, (14) tells us that the center row of P^t contains the t -steps ahead state price approximation. We then solve for the smoothing parameter ζ that minimizes:

$$\min_{\zeta} \quad D_{KL}(S \| U)$$

where D_{KL} denotes the generalized Kullback-Leibler divergence, which is defined as:

$$\sum_{i,j} S_{i,j} \cdot \log \left(\frac{S_{i,j}}{U_{i,j}} \right) - \sum_{i,j} S_{i,j} + \sum_{i,j} U_{i,j} \quad (15)$$

To summarize, as for the estimation of the state price density surface, we generate several solutions that are subsequently evaluated for properties other than the error metrics defining the loss function used to compute them.

3 Empirical Results

Let s denote the spot of the underlying asset and define moneyness as $\psi := x/s$. We estimate state prices over a fixed domain, with $\psi \in [0.4, 1.6]$ and $\tau \in [20, 730]$, using daily closing prices of out-of-the-money (OTM) call and put options on the S&P 500 for each Wednesday between January 5, 2000 and December 26, 2012. The choice of only working with OTM options is motivated by the fact that they are more liquid. In case a particular Wednesday was a holiday, we use the preceding trading day. Option data and interest rates were obtained from OptionMetrics Ivy DB database. In order to determine estimates for all maturities, we linearly interpolate the observable rates. We take the mean of bid and ask quotes as option prices and discard observations with bids below \$0.50, or outside the moneyness-maturity domain of our surfaces. We also exclude options that violate arbitrage constraints. Following Aït-Sahalia and Lo (1998) we use the put-call parity to derive the implied forward from close to at-the-money (ATM) call and put pairs. Given the implied forward, we back out the implied dividend yield via the spot-forward parity and translate OTM puts into in-

the-money (ITM) calls. The resulting data set contains 177'032 call prices spanning a period of 13 years.

Implementation

We use BMS, to map prices to implied volatilities (IVs) and augment the training data with observations at the average ATM IV for $\psi \in [1.2, 1.6]$ and $\tau \in [10, 20]$. This curbs calendar arbitrage by counteracting the effect discrete quotes have on the IVs of short term OTM calls. We also replicate the first options series with $\tau \geq 20$ for $\psi \in [0.9, 1.1]$ at $\tau = 10$. This stabilizes the results over time by reducing the impact of near maturity options dropping out of our model domain. We hasten to add that these points only serve to provide guidance during training and do not affect the ranking of valid solutions. We then map the data into total implied variance space and train several neural networks, of the form in (8). Each network varies both with respect to α and β , which we initialize randomly following the algorithm proposed by Nguyen and Widrow (1990), and the number of basis expansions $H = \max(10, 10 + \lceil \eta \rceil)$, where $\eta \sim \mathcal{N}(0, 2)$. Since options close to ATM are more liquid, we weight the training errors with $\omega = \mathcal{N}(\psi|1, 0.2) + \mathcal{N}(\psi|1, 0.1)$. We train each network for 10 iterations by minimizing the residual sum-of-squares using the method proposed by Foresee and Hagan (1997). Once we have a set of 15 valid surfaces we obtain the final estimator by averaging the 5 solutions with the smallest error on market quotes.

The resulting surfaces span 1201 states and 711 tenors. We use local polynomial smoothing to remove any kinks along the term structure of interest rates and dividend yields to avoid artifacts when mapping

the IV surface to option prices and compute state prices via numerical differentiation. In order to reduce the computational cost of finding the optimal ζ , we downsample S to 150 states.

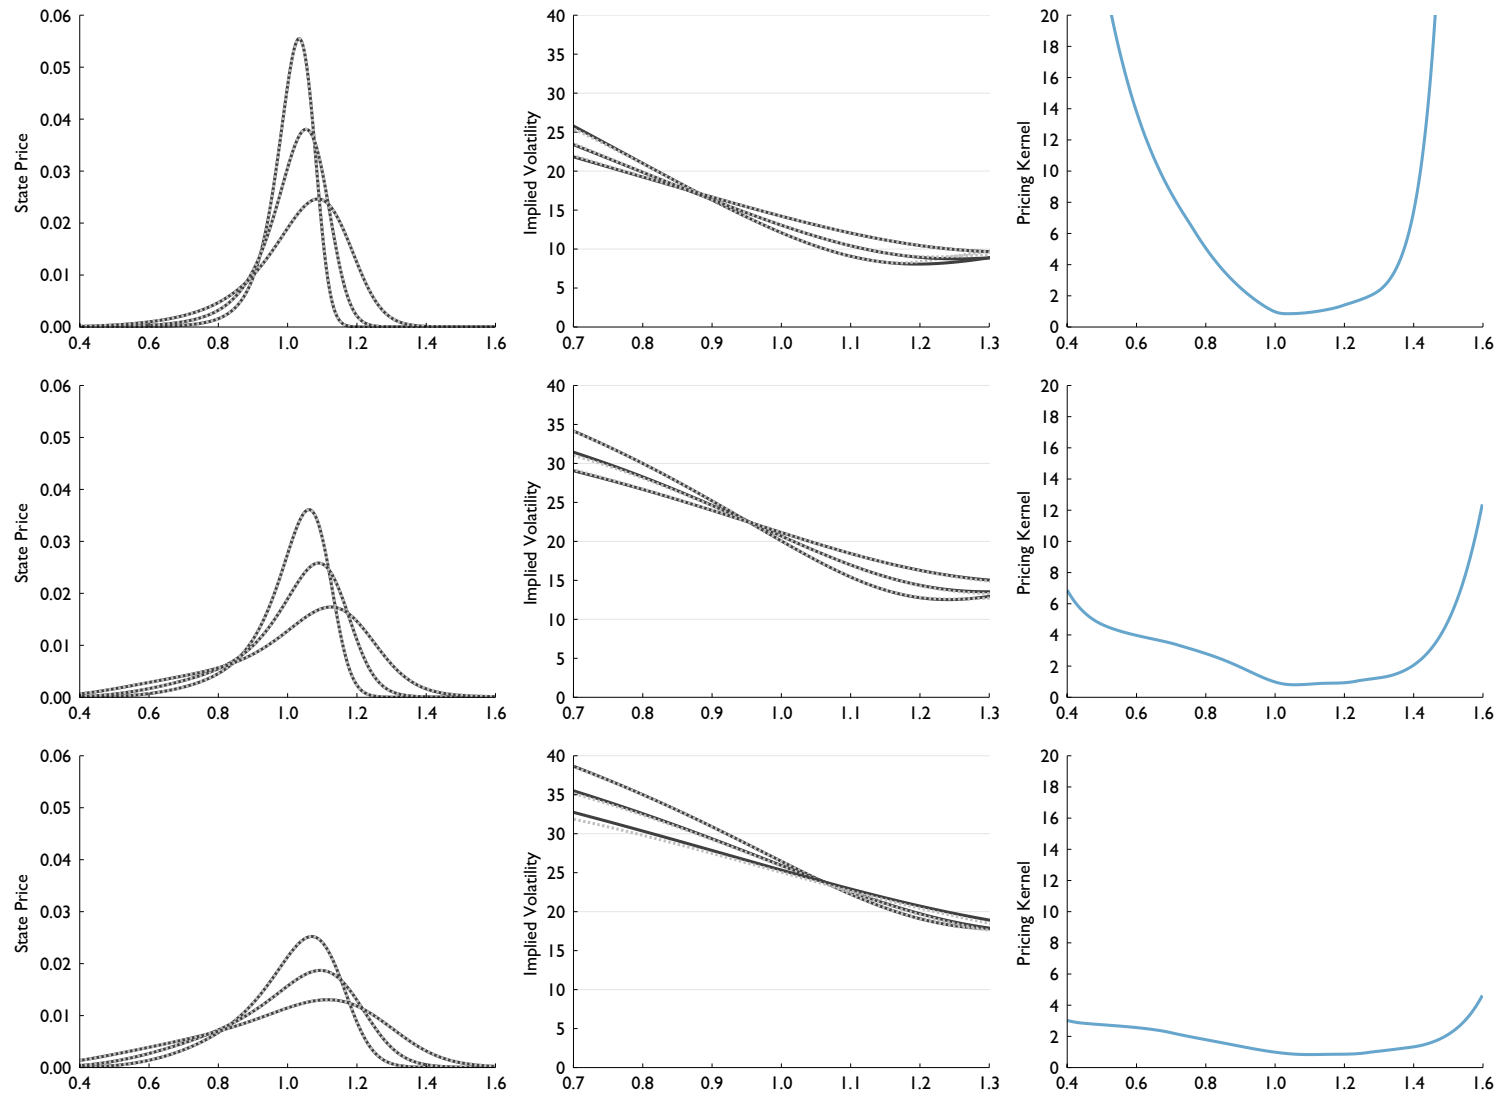
We then slice $S_{150 \times 711}$ with an offset of 90 days to create $A_{621 \times 150}$ and $B_{621 \times 150}$ and using (13), estimate *quarterly* transition matrices for $\zeta \in (0, 0.2]$. In order to prevent a divide by zero error, when computing the generalized Kullback-Leibler divergence of the corresponding Markov chains, we add 10^{-20} to S and U in (15). We choose the transition matrix whose unrolled chain best approximates the state price density surface obtained from the shape constrained neural networks and compute the pricing kernel as the inverse of its right dominant eigenvector. We use cubic spline interpolation to obtain \mathbf{z} for 1199 states and recover the unnormalized real-world probabilities $R_{1199 \times 711}$ from the risk-neutral density via:

$$R_{i,j} = \mathbf{z}_i \cdot S_{i,j} \tag{16}$$

From the risk-neutral and recovered real-world measures we can directly compute moments, that is, we do not need to mimic skewness or kurtosis estimates via the option spanning approach developed in Carr and Madan (2001) and Bakshi, Kapadia, and Madan (2003). Given that our surfaces cover a large and symmetric moneyness range, they are likely less afflicted by the biases documented in Dennis and Mayhew (2002) and Jiang and Tian (2007).

Figure 1 illustrates the quality of our Markov chains by contrasting the initial neural network estimates to the one-, two-, and four-step ahead approximations, both in state price and implied volatility space. The estimates are so close together that they can hardly be distinguished in

Figure 1. State Prices, Implied Volatilities and Pricing Kernels



Notes. The first columns contrast the initial neural network estimates (solid) to those obtained from unrolling the Markov chain (dotted) to 110, 200 and 380 days to maturity. The implied volatilities were obtained by first integrating the Markov chain estimates to option prices. The last column illustrates the corresponding recovered pricing kernels for January 4, 2006, October 24, 2007, and March 19, 2008.

the graph, only in IV space we can discern differences.

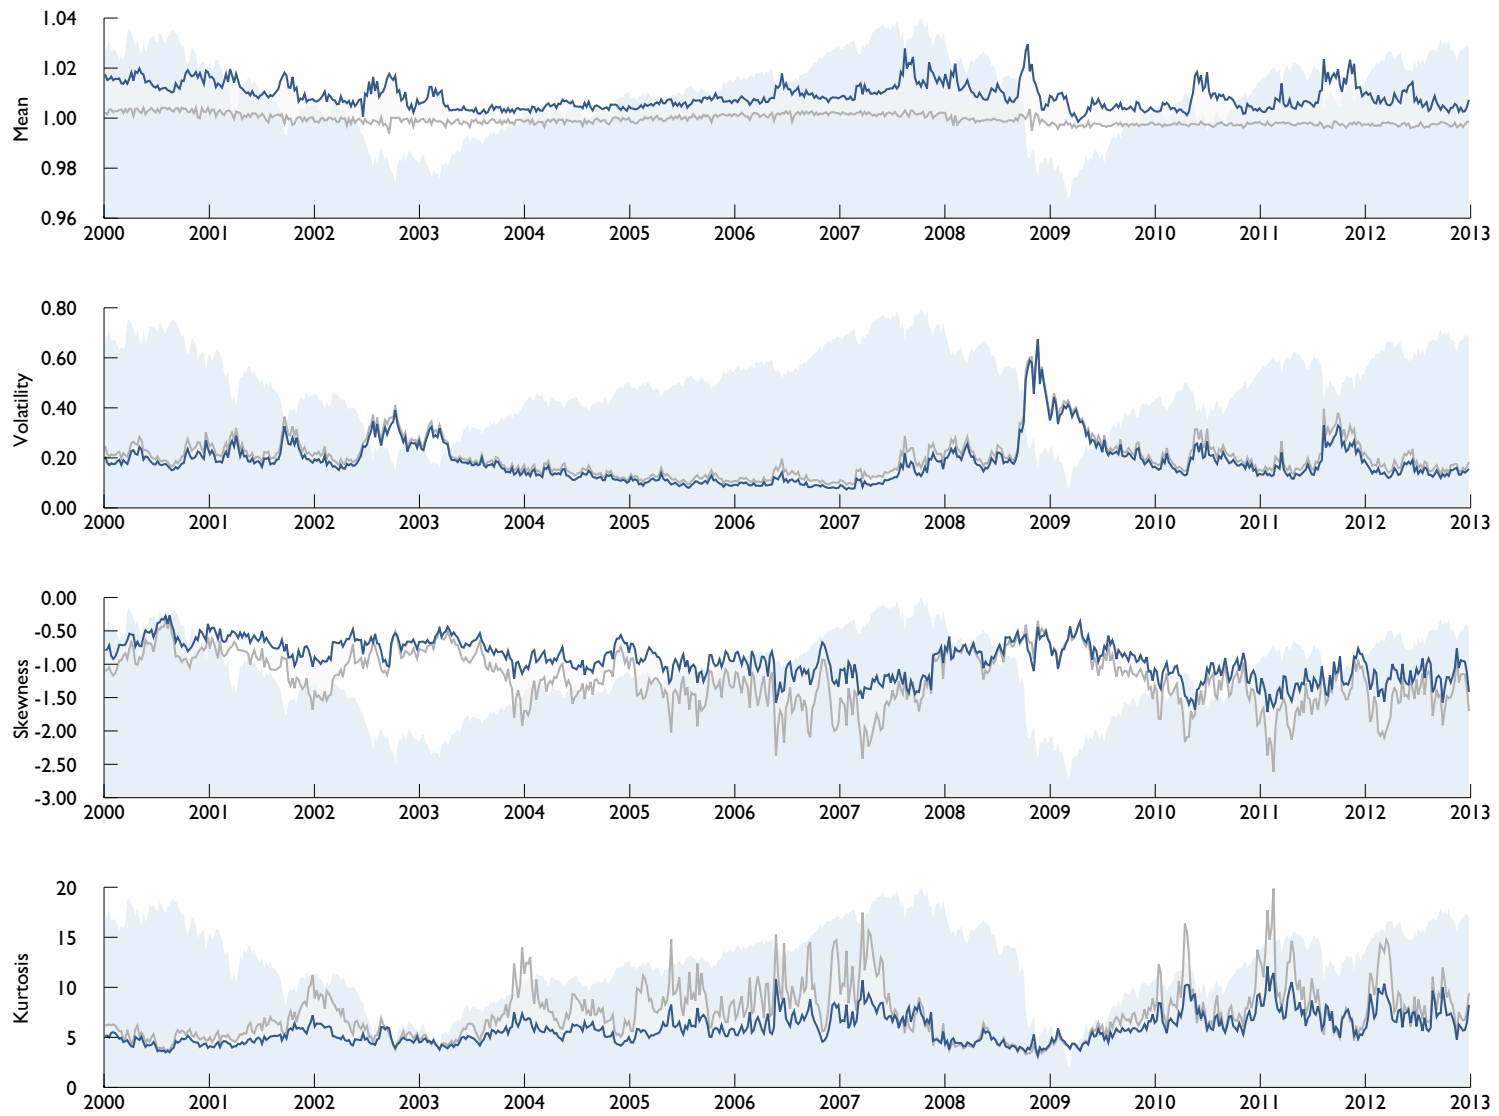
The figure also shows how the increasing uncertainty in the run-up to the crisis of 2008–2009 manifests in the different spaces. We observe how the increasing volatility spreads the state prices over a wider range while pushing the implied volatility curves up. We further see that there is a relation between the state price density and the shape of the corresponding pricing kernel. While such a relation would be expected for the left principal eigenvector \mathbf{v} , it is a priori not obvious for the right Perron vector. We find a strong correlation between state prices averaged across maturities and \mathbf{z} , with an R^2 of 80% over the period from 2000–2012. The relationship is time-varying and weakens during market downturns. This finding sheds light on what drives the recovered pricing kernel and hints at the tantalizing possibility for *pseudo* recovery: sidestepping the construction of the Markov process via $\tilde{\mathbf{z}}_i \propto \sum_j S_{i,j}$.

The recovered pricing kernels are furthermore U-shaped, which is in line with findings of Aït-Sahalia and Lo (2000), Jackwerth (2000), Rosenberg and Engle (2002), Carr et al. (2002), Bakshi, Madan, and Panayotov (2010) and more recently Christoffersen, Heston, and Jacobs (2013).

Moments

Figure 2 depicts the first four moments computed from the 30 day to maturity cross-sections of the risk-neutral and recovered probability densities. Both the range and time-variation of the risk-neutral moments are consistent with results based on the method of Bakshi, Kapadia, and Madan (2003), see e.g., Christoffersen, Jacobs, and Chang (2013). In line with Birru

Figure 2. Risk-Neutral and Recovered Moments



20

Notes. This figure contrasts the evolution of the index (shaded) with that of risk-neutral (light) and recovered moments (dark) for $\tau = 30$.

and Figlewski (2012), we find that skewness becomes less negative leading up to the crisis of 2008–2009, along with a drop in kurtosis.


This pattern persists after the crash, with skewness becoming more negative during the rebound, along with an increase in kurtosis. We further observe the expected correlation between recovered means and volatility, which in turn exhibits the well documented inverse relation to index returns. We remark that the former corroborates the validity of the recovery theorem. While parametric models often explicitly tie the equity risk premium to be proportional to spot variance, see e.g., Filipovic, Gourier, and Mancini (2013), our results are entirely data driven.

We also find that variations in the recovered skewness and kurtosis estimates are considerably more stable over time. This is intuitively appealing, as we would expect agents to be more measured when it comes to updating their views on the real-world distribution, as opposed to the risk-neutral density, which also reflects risk-aversion. This self-stabilizing property of recovery is also consistent with the observed correlation between state prices and \mathbf{z} .

Risk Premia

Figure 3 shows the risk premia for a maturity of 30 days. As expected, the equity risk premium is positive and exhibits significant time-variation with countercyclical behavior, while the variance risk premium is negative throughout and procyclical. These findings are in line with recent studies, e.g., Bollerslev, Tauchen, and Zhou (2009) and Martin (2013). The increase in the equity risk premium during downturns highlights the fact that the average investor requires higher returns to compensate for an increase in

Figure 3. Recovered Risk Premia



22

Notes. This figure contrasts the evolution of the S&P 500 (shaded) with the recovered risk premia for $\tau = 30$ days. With the exception of the variance risk premium, they correspond to the differences between the recovered and risk-neutral moments depicted in Figure 2 above.

return variance. The negative variance risk premium, on the other hand, indicates a preference for lower variance on average to compensate for a higher variance of return variance. As expected, we find that the skewness risk premium is higher during periods where skewness exhibits a higher variance, consistent with the reasoning behind the equity and variance risk premium. For the kurtosis risk premium we observe the same pattern, the higher the variance of kurtosis, the stronger the risk adjustment. The skewness risk premium closely tracks movements in the index, but precedes the sharp drop in 2008. Interestingly, just before the peak of the S&P 500 in 2007, the real-world density exhibits fatter tails than the risk-neutral, something we otherwise only observe just before market rebounds. Furthermore, the skewness premium also changes sign during 2002–2003, 2007–2009, and for the sovereign debt crisis in late 2011.


A Simple Trading Strategy

The relation between moments and the underlying index price motivates us to investigate whether changes in moments predict subsequent returns. Economic intuition dictates that risk-averse investors prefer portfolios with lower volatility, higher skewness, and lower kurtosis and need to be compensated with higher expected returns if any of them increase in modulus, see e.g., Kimball (1993).

We sidestep the well documented pitfalls of predictive regressions⁴ and analyze the viability of market timing. We deliberately choose a simple trading strategy that does not rely on the calibration of any tuning

⁴E.g., Stambaugh (1999), Ferson, Sarkissian, and Simin (2003), Valkanov (2003), Boudoukh, Richardson, and Whitelaw (2008), Welch and Goyal (2008), Anderson, Bianchi, and Goldberg (2013).

Figure 4. Trading at a Moment's Notice



24

Notes. This figure contrasts strategies based on changes in recovered (top panel) and risk-neutral moments (bottom panel) to the index.

parameters. If the absolute value of a moment is larger than it was the week before, we go long the S&P 500 for one week, otherwise we short the index for one week. Figure 4 illustrates the evolution of log-returns for the recovered (top panel) and risk-neutral moments (bottom panel). The bold lines trace the average returns across portfolios.

While the return on the index is essentially zero after 13 years, the yearly returns of the average recovered and risk-neutral strategies are 8.1% and 4.0% respectively.⁵ Strategies based on risk premia fare only slightly better than those based on risk-neutral moments and are not reported. This finding indicates that the combination of changes in level and risk-aversion is more powerful than either signal in isolation. Another notable feature is the stability of the recovered strategies over vastly different regimes and their modest dispersion around the average.⁶

The best performing recovered moment up to the large market drop in the Fall of 2008 is the expected mean return, after that, we see recovered skewness taking over. The risk-neutral moment with the highest performance is variance, while risk-neutral kurtosis is quite inconsistent. The average number of round-trip transactions is 30 per year, which given a conservative estimate of 2 basis-points for commissions and transaction costs, only impacts the yearly returns of our trading strategies by about 0.6%.

Our findings that changes in the modulus of variance, skewness and kurtosis

⁵The fact that such a trading strategy is feasible despite noisy market quotes is a testament to the quality and robustness of the surfaces obtained via the neural network method of Ludwig (2012).

⁶To preempt concerns regarding the impact of the asynchronous recording between option prices and the S&P 500 in OptionMetrics Ivy DB database we implement all strategies using the opening spot of the following trading day, see e.g., Battalio and Schultz (2006).

are positively related to future index returns are consistent with economic theory. The results on risk-neutral skewness and kurtosis are also in line with findings of Conrad, Dittmar, and Ghysels (2013), which is curious because they average the Bakshi, Kapadia, and Madan (2003) moments over three months, thereby diluting the ex-ante nature of the raw estimates, as rightly pointed out by Rehman and Vilkov (2012). In contrast, our risk-neutral and real-world moments are genuinely conditional and forward-looking, as they are constructed from closing market quotes of one single day.

Boyer, Mitton, and Vorkink (2010) and Amaya et al. (2013) find a negative relation between realized skewness and future returns in the cross-section, while Rehman and Vilkov (2012) find the opposite relation using daily risk-neutral skewness estimates obtained via option spanning.

The predictive information content of option implied measures is typically attributed to market fragmentation, with informed agents preferring to express their views on the future via the options market, see e.g., Easley, O'Hara, and Srinivas (1998). The operative word being *informed*. Increased buying pressure for OTM puts would then be indicative of negative future returns. Given that the left tail of our densities is essentially driven by OTM puts that have been converted to ITM calls, higher put prices would likely increase both variance and kurtosis and reduce skewness. We, however, go long the index in this scenario and, over the 676 weeks that we are in the market, make quite consistently money from our trades.

Our results do not support the frequently postulated informational advantage of option traders, but some form of overreaction, which on average is being corrected over the course of the subsequent week. This is in line with findings of Birru and Figlewski (2012), who document overshooting in

RNDs in response to movements in the S&P 500 index. Bates (1991) also reports subsiding fears leading up to the market peak in 1987, and a surge immediately *after* the crash. Lynch and Panigirtzoglou (2008) find that while RNDs change in response to market events, they do not predict them. The fact that trading strategies based on recovered real-world moments perform better than those based on their risk-neutral counterparts suggest that risk-neutralization decreases the signal to noise ratio, and implies that beliefs about properties of the physical process also do not lead the market, but lag behind events that have already unfolded.


We conclude by noting that the performance of market timing based on the recovered mean confirms the validity of the recovery theorem, as the expected real-world return contains much of the same information our contrarian higher-order strategies trade on.

Robustness Checks

Given that the construction of a transition matrix from state prices poses a discrete ill-posed problem, robustness is of paramount importance. If the response to small changes in the state prices is not proportionate and localized, we will not obtain sensible changes in the pricing kernel over time. The robustness of a Markov chain can be quantified by the *eigengap*, that is, the ratio of the dominant to the subdominant eigenvalue. If the value of the subdominant eigenvalue is close to one, the corresponding Markov chain is ill-conditioned in the sense that the stationary distribution is sensitive to perturbations in the transition matrix, cf. Meyer (1994).

The average eigengap of our transition matrices is 0.74. We also find that the

Figure 5. Risk Premia based on *Pseudo* Recovery



28

Notes. This figure illustrates the evolution of risk premia obtained from *pseudo* recovery, that is, a pricing kernel directly derived from S .


social rate of discount δ , that is, the dominant eigenvalue is always smaller than one and exhibits a sensible variation over time. This is as it should be and follows naturally from our algorithm, which only imposes non-negativity and a modicum of smoothness.

Furthermore, we observe that our results are largely robust to changes in the number of states used in the construction of the transition matrices. Changes in the step size do have some impact on the results, which is to be expected, given that the actual risk-neutral process likely violates the assumption of time-homogeneity, see e.g. Carr and Yu (2012). Larger step sizes seem to yield pricing kernels that are more stable across moneyness, which is important when it comes to recovering sensible real-world densities at longer maturities. Last but not least, we find that our results are not sensitive to an expansion of the moneyness domain beyond $\psi \in [0.4, 1.6]$, that is, a larger moneyness domain does not change the shape of the eigenvectors over the original states. This ties back again to the relation between \mathbf{z} and state prices, which are also only locally affected by a wider ψ domain.

Pseudo Recovery

Figure 5 illustrates the risk premia obtained via *pseudo* recovery, that is, the pricing kernel corresponds to the inverse of state prices averaged across maturities. Both the range and time-variation of the risk premia are largely in line with those depicted in Figure 3. The variance risk premium exhibits the most prominent differences with a significantly larger spike in the Fall of 2008, necessitating a change of scale to accommodate it. We also observe lower equity risk premia during times of market turmoil.

Figure 6. Market Timing based on *Pseudo* Recovery



30

Notes. This figure contrasts strategies based on changes in *pseudo* recovered moments (top) and *pseudo* risk premia (bottom) to the index.

Results for the trading strategy are also surprisingly similar. The recovered mean dominates up to the crash in late 2008, while recovered skewness delivers the highest final return. The variance based timing strategy closely tracks the average over all portfolios and exhibits the lowest variation over time. Kurtosis performs the worst.

In contrast to Figure 4, the bottom panel of Figure 6 depicts the results for a strategy based on risk premia. In line with Bollerslev, Tauchen, and Zhou (2009), Goyal and Saretto (2009) and Drechsler and Yaron (2011), we now find an excellent performance for the variance risk premium, closely followed by the equity risk premium. While skewness and kurtosis risk premia perform rather inconsistent, market timing based on the variance risk premium would have yielded a return of over 308% for the period from 2000 to 2013. From a practical perspective, *pseudo* recovery brings with it the additional advantage of computational efficiency.

4 Conclusion

Any empirical implementation of the recovery theorem requires the solution of two ill-posed problems. The first involves estimating the second derivative of the option pricing function from noisy and sparse market quotes to obtain state prices over a fixed domain. The second step entails the construction of a transition matrix that captures the dynamics of said state prices. Only then can recovery be obtained via spectral decomposition. In this paper we present a method based on Tikhonov regularized non-negative least squares to construct robust time-homogeneous Markov chains that not only yield an excellent fit to state prices at various maturity horizons, but pricing

kernels that exhibit a sensible variation over time. We investigate the predictive information content of the risk-neutral and recovered moments and find that – in line with economic theory – the real-world mean impounds most of the information of higher-order risk-neutral moments. Without any restrictions on the shape of the pricing kernel we always find a positive equity risk premium and a negative variance risk premium. Higher-order premia are also consistent with economic intuition. We find that the recovered pricing kernels are closely related to state prices and illustrate the possibility for *pseudo* recovery. While we conclude that the recovery theorem adds economic value, our empirical study is only a first foray into what will likely be a fruitful research area for some time to come. Possible extensions include more comprehensive asset allocation studies along the lines of Kostakis, Panigirtzoglou, and Skiadopoulos (2011) and Jondeau and Rockinger (2012). It would also be interesting to apply our methodology to other markets such as the VIX and investigate if the recovered distributions provide additional information about aggregate risk aversion and economic uncertainty.

References

- Aït-Sahalia, Y. and J. Duarte. 2003. “Nonparametric Option Pricing under Shape Restrictions.” *Journal of Econometrics* 116 (1-2):9–47.
- Aït-Sahalia, Y. and A.W. Lo. 1998. “Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices.” *Journal of Finance* 53 (2):499–547.
- . 2000. “Nonparametric Risk Management and Implied Risk Aversion.” *Journal of Econometrics* 94 (1-2):9–51.
- Amaya, D., P. Christoffersen, K. Jacobs, and A. Vasquez. 2013. “Does Realized Skewness Predict the Cross-Section of Equity Returns?” *Working Paper* .

- Anderson, R.M., S.W. Bianchi, and L.R. Goldberg. 2013. “In Search of a Statistically Valid Volatility Risk Factor.” *Working Paper, University of California at Berkeley* .
- Audrino, F. and D. Colangelo. 2010. “Semi-Parametric Forecasts of the Implied Volatility Surface Using Regression Trees.” *Statistics and Computing* 20 (4):421–434.
- Bakshi, G., N. Kapadia, and D. Madan. 2003. “Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options.” *Review of Financial Studies* 16 (1):101–143.
- Bakshi, G., D. Madan, and G. Panayotov. 2010. “Returns of Claims on the Upside and the Viability of U-Shaped Pricing Kernels.” *Journal of Financial Economics* 97 (1):130–154.
- Banz, R.W. and M.H. Miller. 1978. “Prices for State-Contingent Claims: Some Estimates and Applications.” *Journal of Business* 51 (4):653–72.
- Bates, D.S. 1991. “The Crash of ’87: Was it Expected? The Evidence from Options Markets.” *Journal of Finance* :1009–1044.
- . 1996. “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options.” *Review of Financial Studies* 9 (1):69–107.
- . 2000. “Post-’87 Crash Fears in the S&P 500 Futures Option Market.” *Journal of Econometrics* 94 (1-2):181–238.
- Battalio, R. and P. Schultz. 2006. “Options and the Bubble.” *Journal of Finance* 61 (5):2071–2102.
- Birru, J. and S. Figlewski. 2012. “Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008.” *Journal of Financial Markets* 15 (2):151–180.
- Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. New York: Springer.
- Black, F. and M. Scholes. 1973. “The Pricing of Options and Corporate Liabilities.” *Journal of Political Economy* 81 (3):637–654.

- Bliss, R.R. and N. Panigirtzoglou. 2002. “Testing the Stability of Implied Probability Density Functions.” *Journal of Banking & Finance* 26 (2-3):381–422.
- . 2004. “Option-Implied Risk Aversion Estimates.” *Journal of Finance* 59 (1):407–446.
- Bollerslev, T., G. Tauchen, and H. Zhou. 2009. “Expected Stock Returns and Variance Risk Premia.” *Review of Financial Studies* 22 (11):4463–4492.
- Boudoukh, J., M. Richardson, and R.F. Whitelaw. 2008. “The Myth of Long-Horizon Predictability.” *Review of Financial Studies* 21 (4):1577–1605.
- Boyer, B., T. Mitton, and K. Vorkink. 2010. “Expected Idiosyncratic Skewness.” *Review of Financial Studies* 23 (1):169–202.
- Breeden, D.T. and R.H. Litzenberger. 1978. “Prices of State-Contingent Claims Implicit in Option Prices.” *Journal of Business* 51 (4):621–651.
- Canina, L. and S. Figlewski. 1993. “The Informational Content of Implied Volatility.” *Review of Financial Studies* 6 (3):659–681.
- Carr, P., H. Geman, D.B. Madan, and M. Yor. 2002. “The Fine Structure of Asset Returns: An Empirical Investigation.” *Journal of Business* 75 (2):305–333.
- Carr, P. and D.B. Madan. 2001. “Optimal Positioning in Derivative Securities.” *Quantitative Finance* 1 (1):19–37.
- Carr, P. and J. Yu. 2012. “Risk, Return, and Ross Recovery.” *The Journal of Derivatives* 20 (1):38–59.
- Christoffersen, P., S. Heston, and K. Jacobs. 2013. “Capturing Option Anomalies with a Variance-Dependent Pricing Kernel.” *Review of Financial Studies* 26 (8):1963–2006.
- Christoffersen, P., K. Jacobs, and B.Y. Chang. 2013. “Forecasting with Option Implied Information.” In *Handbook of Economic Forecasting*, vol. 2A, edited by G. Elliott and A. Timmermann. Elsevier.
- Conrad, J., R.F. Dittmar, and E. Ghysels. 2013. “Ex Ante Skewness and Expected Stock Returns.” *Journal of Finance* 68 (1):85–124.

- Cox, J.C. and S.A. Ross. 1976. “The Valuation of Options for Alternative Stochastic Processes.” *Journal of Financial Economics* 3 (1-2):145–166.
- Dennis, P. and S. Mayhew. 2002. “Risk-Neutral Skewness: Evidence from Stock Options.” *Journal of Financial and Quantitative Analysis* 37 (3):471–493.
- Drechsler, I. and A. Yaron. 2011. “What’s Vol Got to Do with it.” *Review of Financial Studies* 24 (1):1–45.
- Dubynskiy, S. and R. Goldstein. 2013. “Recovering Drifts and Preference Parameters from Financial Derivatives.” *Working Paper, University of Minnesota* .
- Dumas, B., J. Fleming, and R.E. Whaley. 1998. “Implied Volatility Functions: Empirical Tests.” *Journal of Finance* 53 (6):2059–2106.
- Easley, D., M. O’Hara, and P.S. Srinivas. 1998. “Option Volume and Stock Prices: Evidence on where Informed Traders Trade.” *Journal of Finance* 53 (2):431–465.
- Fengler, M.R. and L.Y. Hin. 2013. “Semi-Nonparametric Estimation of the Call Price Surface under No-Arbitrage Constraints.” *Working Paper, University of St. Gallen* .
- Ferson, W.E., S. Sarkissian, and T.T. Simin. 2003. “Spurious Regressions in Financial Economics?” *Journal of Finance* 58 (4):1393–1414.
- Figlewski, S. 2010. “Estimating the Implied Risk Neutral Density for the US Market Portfolio.” In *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, edited by T. Bollerslev, J.R. Russel, and M.W. Watson. Oxford University Press, 323–354.
- Filipovic, D., E. Gourier, and L. Mancini. 2013. “Quadratic Variance Swap Models.” *Swiss Finance Institute Research Paper* .
- Foresee, F.D. and M.T. Hagan. 1997. “Gauss-Newton Approximation to Bayesian Learning.” In *Proceedings of the 1997 International Joint Conference on Neural Networks*, vol. 3. IEEE, 1930–1935.

- Goyal, A. and A. Saretto. 2009. “Cross-Section of Option Returns and Volatility.” *Journal of Financial Economics* 94 (2):310–326.
- Hens, T. and C. Reichlin. 2013. “Three Solutions to the Pricing Kernel Puzzle.” *Review of Finance* 17 (3):1065–1098.
- Heston, S.L. 1993. “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.” *Review of Financial Studies* 6 (2):327–43.
- Hornik, K., M. Stinchcombe, and H. White. 1989. “Multilayer Feedforward Networks Are Universal Approximators.” *Neural Networks* 2 (5):359–366.
- Huang, D. and I. Shaliastovich. 2013. “Risk Adjustment and the Temporal Resolution of Uncertainty: Evidence from Options Markets.” *Working Paper, University of Pennsylvania* .
- Hull, J.C. and A. White. 1987. “The Pricing of Options on Assets with Stochastic Volatilities.” *Journal of Finance* 42 (2):281–300.
- Jackwerth, J.C. 2000. “Recovering Risk Aversion from Option Prices and Realized Returns.” *Review of Financial Studies* 13 (2):433–451.
- Jackwerth, J.C. and M. Rubinstein. 1996. “Recovering Probability Distributions from Option Prices.” *Journal of Finance* :1611–1631.
- Jiang, G.J. and Y.S. Tian. 2007. “Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index.” *The Journal of Derivatives* 14 (3):35–60.
- Jondeau, E. and M. Rockinger. 2012. “On the Importance of Time Variability in Higher Moments for Asset Allocation.” *Journal of Financial Econometrics* 10 (1):84–123.
- Kimball, M.S. 1993. “Standard Risk Aversion.” *Econometrica: Journal of the Econometric Society* 61 (3):589–611.
- Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos. 2011. “Market Timing with Option-Implied Distributions: A Forward-Looking Approach.” *Management Science* 57 (7):1231–1249.

- Lawson, C.L. and R.J. Hanson. 1974. *Solving Least Squares Problems*. SIAM.
- Ludwig, M. 2012. “Robust Estimation of Shape Constrained State Price Density Surfaces.” *Working Paper, University of Zurich* .
- Lynch, D. and N. Panigirtzoglou. 2008. “Summary Statistics of Option-Implied Probability Density Functions and Their Properties.” *Bank of England* .
- Martin, I. 2013. “Simple Variance Swaps.” *Working Paper, Stanford University* .
- Martin, I. and S.A. Ross. 2013. “The Long Bond.” *Working Paper, Stanford University and MIT* .
- Merton, R.C. 1973. “Theory of Rational Option Pricing.” *Bell Journal of Economics and Management Science* 4 (1):141–183.
- . 1976. “Option Pricing when Underlying Stock Returns Are Discontinuous.” *Journal of Financial Economics* 3 (1-2):125–144.
- Meyer, C.D. 1994. “Sensitivity of the Stationary Distribution of a Markov Chain.” *SIAM Journal on Matrix Analysis and Applications* 15 (3):715–728.
- . 2000. *Matrix Analysis and Applied Linear Algebra*. SIAM.
- Nguyen, D. and B. Widrow. 1990. “Improving the Learning Speed of 2-Layer Neural Networks by Choosing Initial Values of the Adaptive Weights.” In *Neural Networks, 1990., 1990 IJCNN International Joint Conference on*. IEEE, 21–26.
- Ratcliff, R. 2013. “Relative Option Prices and Risk-Neutral Skew as Predictors of Index Returns.” *The Journal of Derivatives* 21 (2):89–105.
- Rehman, Z. and G. Vilkov. 2012. “Risk-Neutral Skewness: Return Predictability and Its Sources.” *Working Paper, Goethe University Frankfurt* .
- Rosenberg, J.V. and R.F. Engle. 2002. “Empirical Pricing Kernels.” *Journal of Financial Economics* 64 (3):341–372.
- Ross, S.A. 1976. “Options and Efficiency.” *The Quarterly Journal of Economics* 90 (1):75.
- . 2013. “The Recovery Theorem.” *Journal of Finance (forthcoming)* .
- Shimko, D. 1993. “Bounds of Probability.” *Risk* 6 (4):33–37.

- Stambaugh, R.F. 1999. “Predictive Regressions.” *Journal of Financial Economics* 54 (3):375–421.
- Stein, E.M. and J.C. Stein. 1991. “Stock Price Distributions with Stochastic Volatility: An Analytic Approach.” *Review of Financial Studies* 4 (4):727–752.
- Tikhonov, A.N. and V.Y. Arsenin. 1977. *Solutions of Ill-Posed Problems*. V. H. Winston & Sons.
- Valkanov, R. 2003. “Long-Horizon Regressions: Theoretical Results and Applications.” *Journal of Financial Economics* 68 (2):201–232.
- Welch, I. and A. Goyal. 2008. “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction.” *Review of Financial Studies* 21 (4):1455–1508.
- Whaley, R.E. 2000. “The Investor Fear Gauge.” *Journal of Portfolio Management* 26 (3):12–17.
- Xing, Y., X. Zhang, and R. Zhao. 2010. “What Does the Individual Option Volatility Smirk Tell Us about Future Equity Returns?” *Journal of Financial and Quantitative Analysis* 45 (3):641.
- Ziegler, A. 2007. “Why Does Implied Risk Aversion Smile?” *Review of Financial Studies* 20 (3):859–904.