The Ross Recovery Stochastic Discount Factor

Marco Menner *

November 8, 2018

Abstract

This paper explores the recovered stochastic discount factor (SDF) based on Ross (2015) from an empirical perspective. The recovered SDF is not able to price the S&P 500 and the risk-free asset and deviates significantly from alternative SDFs that can price both assets. The theoretical literature showed that the Ross recovery SDF fails to identify a crucial, permanent SDF component and equals the residual transitory SDF component. This paper finds that the empirical Ross recovery SDF does not resemble the empirical transitory SDF component. It also demonstrates the virtual risk-neutrality of recovered SDFs once reasonable economic constraints were added.

Keywords: Ross recovery, stochastic discount factor, transitory component, permanent component, risk-neutral density, transition state prices

*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
I thank Jamie Alcock, Marcel Fischer, Axel Kind, Jens Jackwerth, seminar participants at the University of Konstanz, and seminar participants at the University of St. Gallen for insightful comments.

I. Introduction

The recovery of physical probabilities from available market information is widely regarded as a challenging task. Several recovery methods are based on state prices, which can be obtained from option prices. The stochastic discount factor (SDF), which contains information on the representative investor's preferences, is then used to change a sample of state prices into physical probabilities. Differently, the "pricing kernel puzzle" literature (surveyed in Cuesdeanu and Jackwerth (2018)) starts out with state prices and physical probabilities in order to find the stochastic discount factor. Without making additional assumptions, both physical probabilities and the stochastic discount factor are, however, unknown and cannot be determined simultaneously by reference to state prices alone.

The work of Ross (2015) is based on a given a set of economic assumptions, which led to the formulation of a recovery theorem that can recover both the SDF and physical probabilities from option prices.[1] Jackwerth and Menner (2018) use density tests to demonstrate that the future realized S&P 500 returns do not match such recovered physical probabilities. They find that a most basic implementation of Ross recovery, which uses the minimum assumptions required for applying the theorem, is ill-conditioned and implies unreasonably high state-dependent risk-free rates. The addition of reasonable economic constraints improves numerical stability but pushes recovered physical probabilities towards risk-neutral ones. In this paper, I provide a detailed analysis of the properties of recovered stochastic discount factors.

First, I implement a novel approach to evaluate the performance of the recovery theorem by testing whether recovered one-month SDFs can be used to price the S&P 500 and the risk-free asset. For each recovery method, I obtain a total number of 380 SDFs, one for each month from April 1986 to December 2017.[2] I then determine 380 realizations of those recovered SDFs, where each realization corresponds to one realized future one-month S&P 500 return. I arrive at one time series of monthly SDF realizations for each recovery method. Next, I test whether each SDF time series can (i) price the S&P 500, (ii) price the risk-free asset, and (iii) simultaneously price the S&P 500 and the risk-free asset. I find strong rejections for all Ross recovery methods. A naive method using a Power Utility SDF with risk-aversion parameter $\gamma = 3$ is able to individually, yet not simultaneously, price both assets. I further introduce the Minimum Variance (MinVar) stochastic discount factor, which is defined as the SDF

[1] I explain the assumptions and the mechanism behind the Ross recovery theorem in detail in Section II.

[2] For September 1992 there are no option data available.

with the lowest variance that perfectly prices both assets. There are many other valid SDFs, which can be decomposed into the MinVar SDF and a component that is orthogonal to the assets' returns and has zero expectation, see Cochrane (2000) page 68. Yet, an SDF with a low variance of that residual component can be helpful in reducing the variance of pricing assets. For each recovery method, I estimate the residual component by subtracting the MinVar SDF from the recovered SDF. For Ross Stable and Power Utility, I cannot reject the hypothesis that the residual component has zero expectation. Moreover, the Power Utility SDF has a residual component with a much lower variance than the components of the Ross recovery stochastic discount factors.

In an additional study, I focus on the decomposition of recovered SDFs. The literature on stochastic discount factor decomposition (see Alvarez and Jermann (2005)) allows examining the theoretical properties of the Ross recovery SDF. Prior to the work of Ross (2015), Hansen and Scheinkman (2009) showed that, given the underlying Markovian environment (which is also assumed by Ross), the Perron-Frobenius theorem can be used to recover physical probabilities p. They further argue that those probabilities provide useful long-term insights into risk pricing. One can then relate the probabilities p via a multiplicative adjustment to the spot state prices. I concur with the studies of Alvarez and Jermann (2005) who call this adjustment the transitory stochastic discount factor component.

Further, Hansen and Scheinkman (2009) relate the true physical probabilities via another multiplicative adjustment to the probabilities p. Building on the work of Alvarez and Jermann (2005), I call this adjustment the permanent stochastic discount factor component, the total stochastic discount factor being the product of the permanent and the transitory components. Borovicka et al. (2016) show that Ross (2015) recovers the probabilities p and, by implicitly setting the permanent stochastic discount factor component to one, Ross interprets p as the true physical probabilities. Notwithstanding, is the permanent component truly one? Bakshi et al. (2017) use options on 30-year Treasury bond futures and solve convex minimization problems to find that this component indicates considerable dispersion and does not equal unity, contrary to the implicit assumption made by Ross. My approach differs in that it directly recovers the SDF according to Ross from S&P 500 index options.

By recovering the probabilities p (and not the true physical probabilities), Ross (2015) should allow me to extract the transitory stochastic discount factor component. Alternatively, Bakshi et al. (2017) extract the transitory component of the stochastic discount factor using data on 30-year Treasury bond futures. I obtain a time series of realizations of the SDF from Ross recovery and implement the approach of Bakshi et al. (2017) in order to derive the transitory SDF component. I show that the two approaches generate different outputs, which is further evidence that Ross recovery has empirical limitations.

One explanation for the empirical failure of the Ross recovery theorem is the flatness of recovered SDFs for methods with reasonable economic constraints added, an argument made by Jackwerth and Menner (2018) based on graphical evidence. I use a statistical testing approach to verify this finding. More specifically, I use pseudo-return draws from the risk-neutral cumulative distribution function (cdf) and test with a Berkowitz (2001) test, a Kolmogorov-Smirnov test, and two versions of the Knüppel (2015) test whether those pseudo-returns come from the recovered cdf. The tests reject less often for the constrained Ross recovery methods than for the basic Ross recovery method or the Power Utility method, which confirms that they recover physical probabilities that are close to risk-neutral ones.

The Ross recovery theorem is extensively discussed in the existing literature. From the empirical perspective, Jackwerth and Menner (2018) provide an in-depth analysis of the theorem and use density tests and a prediction analysis of realized moments to verify its failure. Dillschneider and Maurer (2018) and Yao (2018) repeat the density-testing approach for their implementations of the recovery theorem and arrive at a similar conclusion. Earlier empirical work on Ross recovery is Audrino et al. (2015), who develop a profitable trading strategy that is based on recovered moments. They show that this trading strategy performs worse when using risk-neutral moments instead. However, their implementation crucially pre-determines the structure of the recovered SDF, a problem discussed in Jackwerth and Menner (2018). Another empirical paper on the subject is Jensen et al. (2017), who generalize the Ross recovery theorem by adding structure to the SDF and empirically test that model by regressing recovered moments on realized returns and standard deviations.

Accordingly, my paper contributes to the empirical literature about Ross recovery in several ways. As opposed to Bakshi et al. (2017), who use long-term bond market data to show that the permanent SDF component does not equal unity, I directly recover SDF realizations using S&P 500 index options. I am the first to show that such recovered SDFs fail in pricing the market and the risk-free asset. By also finding that recovered SDFs do not match the transitory SDF component, my approach brings together the empirical Ross recovery literature and the theoretical literature on the Ross recovery SDF. I further add to understanding why the Ross recovery theorem fails by using density tests to confirm that constrained Ross recovery methods virtually recover risk-neutrality.

Relevant theoretical papers are Carr and Yu (2012), Qin and Linetsky (2016), Qin et al. (2016), Dubynskiy and Goldstein (2013), and Walden (2017). They mainly analyze implied properties of Ross recovery or focus on extensions of the theorem to continuous time.

The remaining part of this paper is organized as follows. Section II introduces the Ross recovery theorem. Section III explains the methods used to recover physical distributions as well as the statistical tests that are used to test recovered distributions for risk-neutrality. Section IV provides the empirical results of my analyses, including robustness. In Section V, recovered distributions are tested for risk-neutrality. Section VI concludes.

II. Ross Recovery

The Ross recovery theorem relies on certain economic assumptions and so-called *transition state prices*. Transition state prices differ from regular *spot* state prices: spot state prices $\pi_{0,j}$ are nothing but discounted risk-neutral probabilities of moving from the current state 0 to another state j and can be easily computed from option prices using Breeden and Litzenberger (1978). Transition state prices $\pi_{i,j}$, on the other side, are discounted risk-neutral probabilities of moving from *any* given state i to another state j. Thus, more information is needed to determine such state prices, which commence at initial states i that do not necessarily equal the current state 0 of the economy.

The trick for recovery lies in the three assumptions underpinning the theorem. First, Ross (2015) requires transition state prices $\pi_{i,j}$ to be positive, which is a reasonable assumption. Second, transition state prices need to be time-homogeneous. This assumption, although strong, is crucial to achieving recovery, since it allows us to back out transition state prices $\pi_{i,j}$, which have to be determined for *several* initial states and *one* maturity (this study relies on a maturity of one month) by using available spot state prices $\pi^t_{0,j}$ for *one* initial state (the current state of the economy) and *several* maturities t. Third, the Ross recovery stochastic discount factor $SDF_{i,j}$ is transition independent, meaning it can be expressed in the form

$$SDF_{i,j} = \delta \frac{u'_j}{u'_i} \tag{1}$$

with a positive scalar δ, which can be interpreted as a utility discount factor, and positive terms u'_j and u'_i, which can be interpreted as state-dependent marginal utilities. Assuming that we have N different states and a known transition state price matrix Π with entries $\pi_{i,j}$, for $i, j = 1, ..., N$, Ross then is able to state the eigenvalue problem

$$\Pi z = \delta z, \quad \text{where} \quad z_i = \frac{1}{u'_i}, \quad i = 1, ..., N, \tag{2}$$

which delivers a unique positive eigenvector z and a positive eigenvalue δ.[3] Physical transition probabilities $p_{i,j}$ are then determined as the ratio between transition state prices $\pi_{i,j}$ and the (transition) stochastic discount factor $SDF_{i,j}$.

III. Methodology

In this section, I explain how I obtain spot state prices from observed option prices and then focus on the extraction of transition state prices from spot state prices. I also implement a method that applies Ross recovery without using transition state prices. Finally, I introduce the Berkowitz test, the Knüppel (2015) test, and the Kolmogorov-Smirnov test, which are applied in a simulation study to test recovered distributions for risk-neutrality.

A. Data

The application of Ross recovery requires me to first calculate spot state prices $\pi^t_{0,j}$ for different states j and maturities t. The current state is labeled with $i = 0$ and comes from a set of state indices I with $I = \{-N_{low}, ..., 0, ..., N_{high}\}$ where $N = N_{low} + N_{high} + 1$. The transition state j is also part of that set I.

Spot state prices are calculated from a rich option price surface. First, I collect prices of European options (computed as averages of bid- and ask quotes) on the S&P 500 with a maturity of up to one year from the Berkeley Options Database (April 1986 - December 1995) and OptionMetrics (January 1996 - December 2017). My sample uses all dates that are 30 calendar days prior the third Friday of each month. Next, I apply standard filters to remove options that violate no-arbitrage conditions. I exclude in-the-money options since they are less liquid than out-of-the-money options. Option prices are then converted into observed implied volatilities.

B. Obtaining Spot State Prices from Observed Option Prices

I follow Jackwerth and Menner (2018) to obtain a rich implied volatility surface from observed implied volatilities. Their approach balances between the sum of squared local total second derivatives of implied volatilities $\sigma''_{i,t}$ (to achieve smoothness of the volatility surface) and the sum of squared differences between model- and observed implied volatilities (to achieve

[3]The uniqueness of a positive eigenvector z and a positive eigenvalue δ is guaranteed by the Perron-Frobenius theorem, which can be applied when the assumptions of the Ross recovery theorem hold.

a good fit of the surface). To balance between smoothness and fit, a trade-off parameter λ is used. The squared second derivatives $(\sigma''_{i,t})^2$ are multiplied with maturity t to give less weight to the short-term volatility smile, which usually has a stronger curvature. I solve

$$\min_{\sigma_{i,t}} \quad \frac{1}{TN_{fine}} \cdot \sum_{t=1}^{T} \sum_{i \in I_{fine}} (\sigma''_{i,t})^2 \cdot t \quad + \quad \lambda \cdot \frac{1}{L} \cdot \sum_{l=1}^{L} \left(\sigma_{i(l),t(l)} - \sigma^{\text{obs}}_{i(l),t(l)}\right)^2 \qquad s.t. \quad \sigma_{i,t} \;\geq\; 0, \tag{3}$$

with $\sigma^{\text{obs}}_{i(l),t(l)}$ being the $l-th$ observed implied volatility (out of L observations) and I_{fine} being a fine set of indexes for states i with a total number of N_{fine} states. The numerical second derivatives $\sigma''_{i,t}$ are computed with a finite differences approach.

For the fine state space, I discretize the strike prices with a step size of \$1.25 and divide strike prices by the level of the S&P 500 index to convert them into moneyness levels. The maturity is discretized with ten steps per month using a total number of 120 steps (i.e., a maximum maturity of one year). All observed options lie on this fine grid. More details on the procedure are given in Jackwerth and Menner (2018). The identification of transition state prices requires an upper bound of the state space dimension, which is set by the maturity dimension of the volatility surface, see Section III.C for details. After solving the optimization in (3), I thus linearly interpolate the fine volatility surface to obtain a coarser volatility surface that is required to compute transition state prices. Further, I ensure that the current state $i = 0$ lies on the coarse volatility surface. Converting this coarse implied volatility surface into option prices and applying Breeden and Litzenberger (1978) finally delivers the spot state prices $\pi^t_{0,j}$ for states $j \in I$ and maturities $t = 1, ..., 120$ that are needed to compute transition state prices.

C. Empirical Ross Recovery: Obtaining Physical Probabilities

Next, I use the methods introduced in Jackwerth and Menner (2018) to extract transition state prices $\pi_{i,j}$ from spot state prices $\pi^t_{0,j}$. Since transition state prices are time-homogeneous, they can be linked to spot state prices by the law of total probability

$$\pi^{t+10}_{0,j} = \sum_{i \in I} \pi^t_{0,i} \cdot \pi_{i,j}, \quad \forall j \in I, \quad t = 0, ..., 110, \tag{4}$$

where $\pi_{0,i}^{t}$ are spot state prices with a maturity indexed by t, $\pi_{i,j}$ are one-month transition state prices, and $\pi_{0,i}^{t+10}$ are spot state prices with a maturity indexed by $t+10$. Because of to the chosen maturity step size of 0.1 months, the maturity indexed by $t+10$ exceeds the maturity indexed by t by exactly one month, which matches the length of one transition period. For $t=0$, spot state prices $\pi_{0,i}^{t}$ are defined as a vector with an entry of one for the current state $i=0$, and entries of zero, otherwise.

Equation 4 links spot state prices with a one-month maturity difference via transition state prices by using an *overlapping* approach. I call this an overlapping approach because this linkage is performed for each point on the fine 0.1-month maturity grid. A *non-overlapping* approach would still link spot state prices with a one-month maturity difference via transition state prices but would do so on a much coarser one-month maturity grid. This, however, would drastically limit the state space dimension: For the non-overlapping approach, one could only use a total number of 13 monthly maturities (including a starting maturity of zero). This means that identification in Equation 4 can only be achieved when limiting the number of states in the set I to N=12. Having only 12 states available, however, leads to extremely coarse spot state prices and transition state prices.

Jackwerth and Menner (2018) thus implement the overlapping approach, which allows to extract one-month transition state prices that are defined on a finer grid with N=111 states. They first introduce a basic Ross recovery method (labeled as **Ross Basic**), which only requires transition state prices to be positive $\pi_{0,j}$. They find that this method delivers jagged transition state prices with many extreme local peaks. Further, rowsums of the transition state price matrix Π are not bounded. This proves to be problematic since transition state prices should, for a given initial state i, sum to that state's risk-free discount factor. Ross Basic thus leads to unreasonably high and low state-dependent risk-free rates. To stabilise transition state prices, Jackwerth and Menner (2018) introduce another method, **Ross Bounded**. This additionally demands all rowsums of Π to lie between 0.9 and 1, which restricts the monthly risk-free rates to between 0% and 11.11%, and delivers smooth transition state prices. Yet another problem is that transition state prices should be higher when the initial state i is closer to the ending state j since it is more likely to end up in a neighboring state than in some far away state. This is neither guaranteed nor observed with the previous methods. In addition to bounding the rowsums, Jackwerth and Menner (2018) thus force the rows in the Π matrix to be unimodal, with the highest values lying on the main diagonal of Π. They label this method **Ross Unimodal**. I implement all the above methods and extract transition state prices with the following least squares approach:

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

$$\textbf{Ross Bounded:} \quad above\ \ and \quad 0.9 \le \sum_{j\in I}\pi_{i,j} \le 1.0 \quad \forall i \in I \tag{5}$$

$$\textbf{Ross Unimodal:} \quad above\ \ and \quad \pi_{i,j} \le \pi_{i,l} \quad \forall j,l,i \in I \quad with \quad j < l \le i$$

$$and \quad \pi_{i,j} \ge \pi_{i,l} \quad \forall j,l,i \in I \quad with \quad j > l \ge i.$$

Next, I apply the Ross recovery theorem to obtain one-month physical transition probabilities $p_{i,j}$ from $\pi_{i,j}$ for each recovery method. For my analysis, I only require the recovered probabilities $p_{0,j}$ for which the initial state i equals the current state 0.

Another method in Jackwerth and Menner (2018), labeled as **Ross Stable**, is able to achieve recovery without using transition state prices altogether. That method exploits the stochastic discount factor structure implied by the recovery theorem. I introduce the idea by considering a t-month transition state price matrix $\Pi^{(t)}$. Due to the time-homogeneity assumption, $\Pi^{(t)}$ can be obtained by multiplying the one-month transition state price matrix Π t-times with itself. Using the eigenvalue problem from Equation 2 gives:

$$\Pi^{(t)}z = \Pi^{(t-1)}\cdot\Pi\cdot z = \Pi^{(t-1)}\cdot\delta\cdot z = \delta\cdot\Pi^{(t-2)}\cdot\Pi\cdot z = ... = \delta^t z \tag{6}$$

Next, I only consider the row of $\Pi^{(t)}$ with current initial state $i = 0$. In theory, that row consists of known spot state prices $\pi_{0,j}^t$. I then make use of Equation 6 to set up the system:

$$\sum_{j\in I}\pi_{0,j}^t\cdot\left[\frac{z_j}{z_0}\right] = \delta^t, \quad t = 1, ..., 120 \tag{7}$$

Directly solving that equation for $\left[\frac{z_j}{z_0}\right]$ does not guarantee the positivity of stochastic discount factors and the utility discount factor δ. I thus follow Jackwerth and Menner (2018) and apply the following least squares approach:

$$\min_{\left[\frac{z_j}{z_0}\right],\delta} \quad \sum_{t=1}^{T}\left(\sum_{j\in I}\pi_{0,j}^t\cdot\left[\frac{z_j}{z_0}\right] - \delta^t\right)^2 \quad s.t. \quad \left[\frac{z_j}{z_0}\right] > 0,\ 1 > \delta > 0. \tag{8}$$

I solve for $\left[\frac{z_j}{z_0}\right]$ and δ, which are used to compute the stochastic discount factor and to transform spot state prices into physical probabilities, see Jackwerth and Menner (2018). The physical probabilities for this method are recovered on N=120 states.

Lastly, I implement a benchmark method that assumes a power utility SDF with risk-aversion parameter $\gamma = 3$. I label that method **Power Utility**. For comparison, physical probabilities are computed on the same N=120 states that are used for Ross Stable.

For all five methods (Ross Basic, Ross Bounded, Ross Unimodal, Ross Stable, and Power Utility), I divide one-month spot state prices by the recovered probabilities to determine the SDFs.[4] For each method, I end up with one SDF for each of the 380 sample days.

D. Density Tests

Jackwerth and Menner (2018) provide graphical evidence that recovered SDFs often tend to be flat, which implies equality of recovered physical and risk-neutral distributions. I introduce four different tests to check whether recovered physical distributions do indeed match the risk-neutral ones. For each month τ in my sample, I first draw one pseudo risk-neutral return r_τ from the risk-neutral cdf Q_τ. Additionally, I recover the physical probability distribution p_τ and the corresponding cdf P_τ for that day. Given $Q_\tau = P_\tau$, the following transformation of $1 + r_\tau$ into u_τ is i.i.d uniformly distributed, $u_\tau \sim i.i.d.\ U(0, 1)$:

$$u_\tau = P_\tau(1 + r_\tau) = \int_{-\infty}^{1+r_\tau} p_\tau(x)dx \tag{9}$$

I generate 10.000 sequences of u_τ (each sequence consisting of 380 uniform points, one for each month) and test whether each sequence is i.i.d. uniformly distributed. I repeat that approach for all recovery methods and report how often each of the following statistical tests rejects uniformity at the 5% significance level.

First, I use the Berkowitz (2001) test, which jointly tests u_τ for uniformity and the i.i.d. property. This test first transforms the series u_τ into another series z_τ with the inverse

[4] Jackwerth and Menner (2018) label this SDF the *implied* SDF. Additionally, they define the *model* SDF $m_{0,j}$ as the ratio between transition state prices $\pi_{0,j}$ and recovered probabilities $p_{0,j}$. Since extracted one-month transition state prices $\pi_{0,j}$ only match one-month spot state price in theory yet not empirically, the model SDF and the implied SDF do not coincide. The main study in this paper uses the implied SDF.

standard normal cdf Φ:

$$z_\tau = \Phi^{-1}(u_\tau) = \Phi^{-1}\left(\int_{-\infty}^{1+r_\tau} p_\tau(x)dx\right) \tag{10}$$

Given $Q_\tau = P_\tau$, it follows that $z_\tau \sim i.i.d.\ N(0,1)$. Berkowitz tests independence and standard normality of z_τ by estimating the following AR(1) model:

$$z_\tau - \mu = \rho(z_{\tau-1} - \mu) + \epsilon_\tau \tag{11}$$

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

$$LR_3 = -2(LL(0,1,0) - LL(\hat{\mu}, \hat{\sigma}, \hat{\rho})) \tag{12}$$

with LL being the log likelihood of Equation 11.

Next, I use a test introduced in Knüppel (2015). Instead of testing u_τ for uniformity, this test first linearly transforms the series u_τ into $y_\tau = \sqrt{12}\,(u_\tau - 0.5)$ and tests y_τ for standardized uniformity with mean=0 and variance=1. Next, the test compares the first S moments of y_τ to the respective moment of its theoretical counterpart in a GMM-type procedure with test statistic α_S:

$$\alpha_S = T \cdot D_S^\intercal \Omega_S^{-1} D_S \tag{13}$$

with D_S being a vector that consists of the differences between sample moments $\frac{1}{\mathcal{T}} \sum_{\tau=1}^{\mathcal{T}} y_\tau^s$ and theoretical moments μ_s for $s = 1, ..., S$. Ω_S is a consistent estimator of the covariance matrix of moment differences between y_τ and its theoretical counterpart. As in Knüppel (2015), I set all covariances in Ω_S between odd and even moment differences to zero. In my analysis, I perform the test by (i) using three moments (S=3) and (ii) using four moments (S=4). The test accounts for serial correlation of u_τ by estimating a Newey-West covariance matrix with automated lag length as suggested by Andrews (1991).

Lastly, I apply a Kolmogorov - Smirnov (KS) test. This test determines the maximum distance between the uniform cdf U and the empirical cdf U_Q of the percentiles u_τ and is based on the following test statistic:

$$KS = \sup_v |U_Q(v) - U(v)| \tag{14}$$

IV. Empirical Results

Recovering wrong physical probabilities is equivalent to recovering the wrong stochastic discount factor. Thus, the empirical usefulness of the Ross recovery theorem stands and falls with the pricing abilities of the recovered SDFs. By switching to a time series perspective, I investigate whether recovered stochastic discount factors are able to price the market and the risk-free asset. I further compare realizations of the recovered stochastic discount factor to theoretically motivated stochastic discount factors, namely, the Minimum Variance stochastic discount factor and the transitory stochastic discount factor component.

A. Do realized SDFs price the S&P 500 and the risk-free asset?

Jackwerth and Menner (2018) show that for a particular day (February 17, 2010), the recovered stochastic discount factors are often rather flat. Yet, how do they behave in general? A time series perspective can help here and I ask for a time series of positive stochastic discount factor realizations m_τ being able to price the market and the risk-free asset.[5] Out of the many stochastic discount factors satisfying these constraints, I introduce the Minimum Variance (MinVar) stochastic discount factor:

$$\min_{SDF_\tau} Var(SDF) \text{ s.t. } \frac{1}{\mathcal{T}}\sum_{\tau=1}^{\mathcal{T}} SDF_\tau \cdot R_\tau = 1, \quad \frac{1}{\mathcal{T}}\sum_{\tau=1}^{\mathcal{T}} SDF_\tau \cdot Rf_\tau = 1, \quad SDF_\tau > 0 \tag{15}$$

with R_τ being the realized 30-day future return market return and Rf_τ the realized 30-day future risk-free return at date τ. Additionally, I compute one time series of realizations of the stochastic discounts for each recovery method. I determine this stochastic discount factor realization for each date τ as the value of the method's implied SDF that corresponds to the market return R_τ.

Figure 1 depicts the realized stochastic discount factors (gray squares) for the four Ross recovery methods (Panel A-D), for Power Utility (Panel E), and for the Minimum Variance stochastic discount factor (Panel F). The black line in each panel represents a Gaussian kernel regression line through the realized SDF values.

The realized stochastic discount factor for Ross Basic moves around widely and demonstrates the instability of that method over time. In comparison, Ross Bounded, Ross Uni-

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

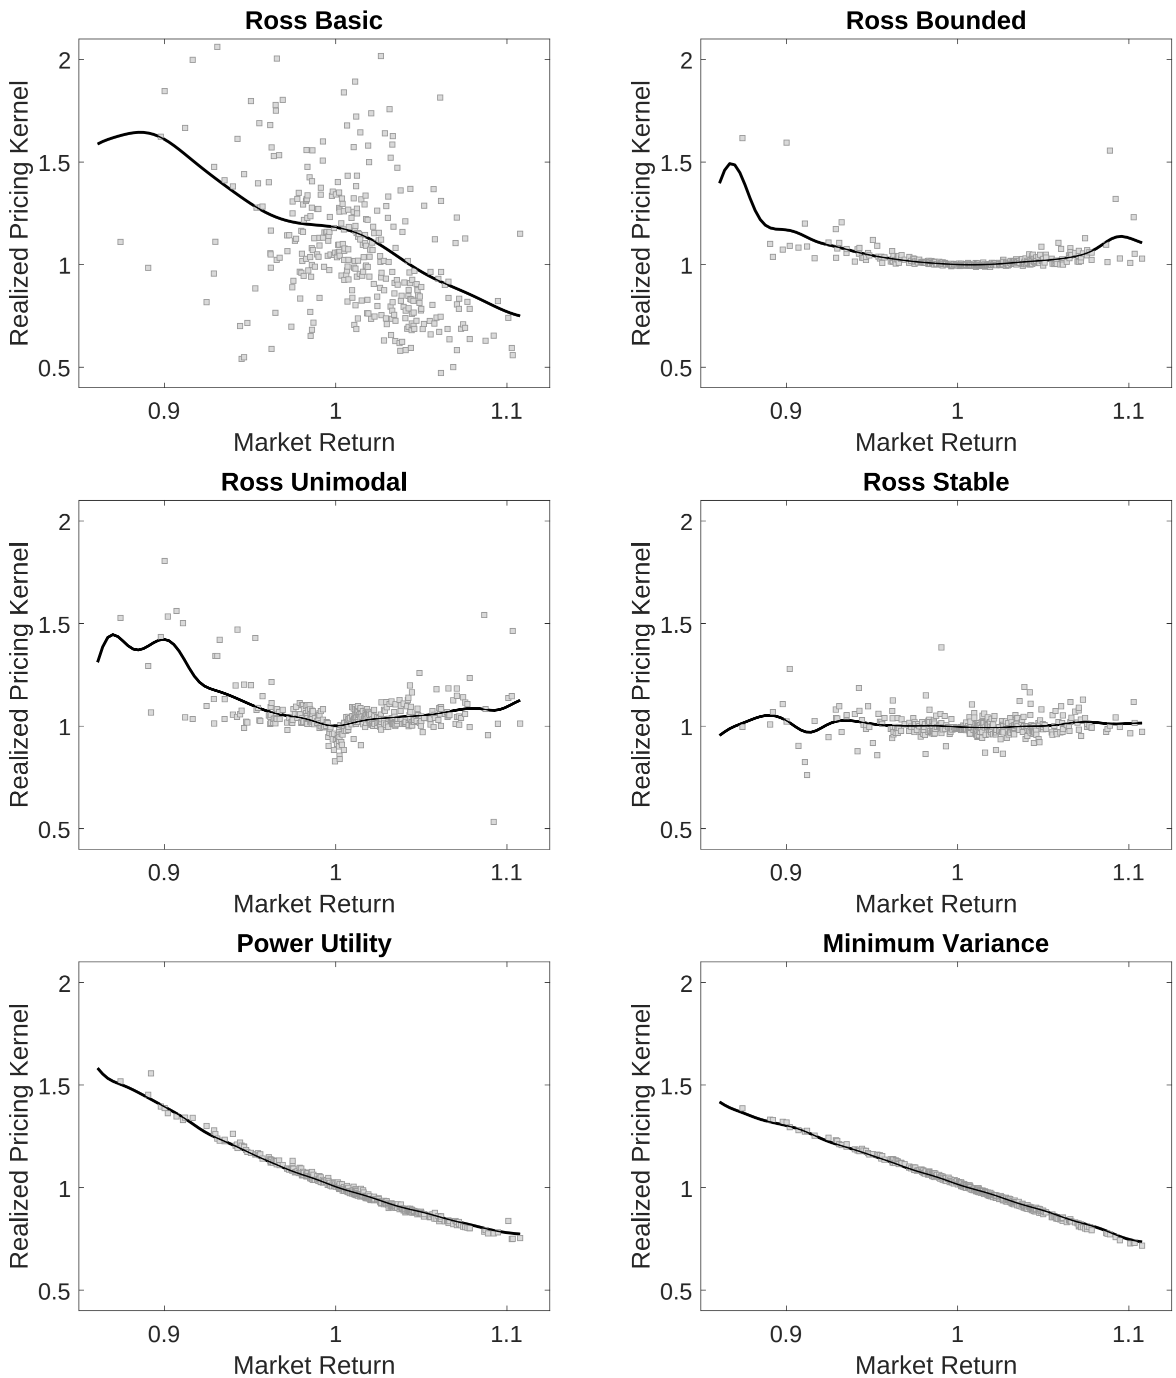

Figure 1. Time Series of Stochastic Discount Factors. I depict a time series of realized stochastic discount factors (gray squares) for different recovery methods. I further depict a Gaussian kernel regression line (black) through the realized stochastic discount factor values. Panel A shows realized stochastic discount factors for Ross Basic, Panel B for Ross Bounded, Panel C for Unimodal, Panel D for Ross Stable, Panel E for Power Utility with $\gamma = 3$, and Panel F for the Minimum Variance stochastic discount factor.

modal, and Ross Stable are less variable and exhibit rather flat stochastic discount factors. These cluster around one, approximating risk-neutrality. The picture is different for Power Utility and the Minimum Variance stochastic discount factor, which are both decreasing in future realized returns.

The stochastic discount factor realizations for the Ross recovery methods and for Power Utility are not constrained to exactly price the S&P 500 and the risk-free asset. This is in contrast to the Minimum Variance stochastic discount factor from Equation 15 and enables a way of testing the pricing ability of recovered stochastic discount factors and, by that, the performance of the Ross recovery theorem. For each recovery method, I test the hypotheses that recovered realized SDFs can ($H1$) price the S&P 500, ($H2$) price the risk-free asset, ($H3$) price the S&P 500 and the risk-free asset simultaneously:

$$H1: \quad E(SDF \cdot R) = \frac{1}{\mathcal{T}} \sum_{\tau=1}^{\mathcal{T}} SDF_{\tau} \cdot R_{\tau} = 1 \tag{16}$$

$$H2: \quad E(SDF \cdot R_f) = \frac{1}{\mathcal{T}} \sum_{\tau=1}^{\mathcal{T}} SDF_{\tau} \cdot Rf_{\tau} = 1 \tag{17}$$

$$H3: \quad E(SDF \cdot R) = 1 \quad and \quad E(SDF \cdot Rf) = 1 \tag{18}$$

Results are given in Table I.

The Minimum Variance SDF perfectly prices both assets. Ross Basic performs worst in pricing the market and the risk-free asset with expected SDF-adjusted returns above 1.3 instead of 1. The other Ross recovery methods perform slightly better with expected values closer to (yet still significantly above) unity. For all Ross recovery methods I can, however, reject the ability to price the S&P 500, the risk-free asset, or both simultaneously. Only for the ability of Ross Stable to price the risk-free asset, I do not find a rejection at the 5% significance level. The Power Utility SDF performs much better. Here, I cannot reject its ability to price both assets individually. However, I can also reject its ability to price both assets simultaneously at the 5% level (yet not at the 1% level).

For robustness, I repeat the above analysis but apply Ross recovery on a reduced state space. While the original analysis uses 120 maturity steps (10 steps per month), the reduced state space only uses 96 maturity steps (8 steps per month). This limits the state space from 111 to 89 states for the methods Ross Basic, Ross Bounded, and Ross Unimodal, and from 120 to 96 states for the methods Ross Stable and Power Utility.

For the sake of robustness, I repeat the original analysis but exclude recession periods classified by the NBER (see e.g. https://fred.stlouisfed.org/series/VIXCLS) from my main sample. Removed periods are the Early 1990s Recession (July 1990 until March 1991), the Early 2000s Recession (March 2001 until November 2001), and the Great Recession (December 2007 until June 2009). This reduces my main sample from 380 dates to 343 dates. Both robustness studies do not qualitatively change my main findings. Result tables are given in Appendix A.1.

Table I. **Test if SDFs price the S&P 500 and the risk-free asset.** I test whether the time series of realized recovered SDFs is able to price the S&P 500 and the risk-free asset. The first column shows the values of $E\left[SDF \cdot R\right]$ for S&P 500 returns R and the p-values for testing this expectation against unity. The second column shows the values of $E\left[SDF \cdot R_f\right]$ for risk-free returns R_f and the again p-values for testing this expectation against unity. The third column shows the p-values of a Wald test that simultaneously tests both expectation values against unity.

SDF from Method	$E\left[SDF \cdot R\right]$ (p-value - unity)	$E\left[SDF \cdot R_f\right]$ (p-value - unity)	**Joint Wald Test** (p-value - joint unity)
Ross Basic $\pi_{i,j} > 0$	1.304 (0.000)	1.308 (0.000)	- (0.001)
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.035 (0.000)	1.030 (0.000)	- (0.001)
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.060 (0.000)	1.056 (0.000)	- (0.000)
Ross Stable No transition state prices	1.010 (0.002)	1.005 (0.076)	- (0.007)
Power Utility with $\gamma = 3$	0.999 (0.895)	1.000 (0.980)	- (0.022)
MinVar	1.000 (1.000)	1.000 (1.000)	- (1.000)

B. How do recovered SDFs compare to the MinVar SDF?

The Minimum Variance SDF is not the only stochastic discount factor that is able price the market and the risk-free asset. Any other valid stochastic discount factor $SDF^{\star}$ can be decomposed into the Minimum Variance stochastic discount factor SDF^{MinVar} and a residual component $\epsilon^{\star}$ that is orthogonal to the returns of the market R and the returns of the risk-free asset R_f:[6]

$$SDF_{\tau}^{\star} = SDF_{\tau}^{\text{MinVar}} + \epsilon_{\tau}^{\star}, \quad \text{where} \quad E(R \cdot \epsilon^{\star}) = E(R_f \cdot \epsilon^{\star}) = 0. \tag{19}$$

Since all valid stochastic discount factors $SDF^{\star}$ have the same expected value, the expected orthogonal component $E(\epsilon^{\star})$ further must equal zero. Moreover, given that the Minimum Variance stochastic discount factor SDF^{MinVar} is uncorrelated with the orthogonal component $\epsilon^{\star}$, a lower variance of $\epsilon^{\star}$ can lead to a lower asset pricing variance and implies an SDF that is more similar to the MinVar SDF.[7] Out of all valid stochastic discount factors, an SDF which comes with a low variance of $\epsilon^{\star}$ would thus be preferable to an SDF which comes with a high variance of $\epsilon^{\star}$.

I am interested in the properties of the empirical residual component $\hat{\epsilon}$ for a given recovered stochastic discount factor SDF_{τ}, which I estimate as:

$$\hat{\epsilon}_{\tau} = SDF_{\tau} - SDF_{\tau}^{\text{MinVar}}. \tag{20}$$

First, I use a t-test to test the hypothesis that $E(\hat{\epsilon}) = 0$, a condition that must be fulfilled for valid stochastic discount factors. Further, I compute the variance of the residual component, $Var(\hat{\epsilon})$. Results are given in Table II. The second column reports the expected value of the residual component $\hat{\epsilon}$ and the p-value for testing whether the residual component $\hat{\epsilon}$ has an expectation of zero. I can reject that hypothesis for all recovery methods but Ross Stable (p-value of 0.45) and Power Utility (p-value of 0.92) at the 5% significance level. This finding shows again that Ross Basic, Ross Bounded, and Ross Unimodal all deliver invalid stochastic discount factors. The variance of the residual component for the Ross recovery SDFs is much higher (values above 0.016 topped by a value of 2.602 for Ross Basic) than for the Power Utility SDF (a value of 0.001).[8] Unreported results further show that the Ross recovery

[6] See Cochrane (2000) page 68 for details on that decomposition.

[7] See Appendix A.4 for details.

[8] For comparison, the variance of the MinVar SDF equals 0.016.

SDFs are highly (negatively) correlated with their residual component. The Power Utility SDF, on the other hand, does not correlate with its residual component.

Ross Basic is so noisy that it hardly bears any resemblance with the Minimum Variance stochastic discount factor. On the other hand, even mild economic constraints in Ross Bounded, Ross Unimodal, or Ross Stable lead to stochastic discount factors that are almost risk-neutral and, again, do not resemble the Minimum Variance stochastic discount factor.

Table II. Analyzing the residual SDF component of recovered SDFs. I present the results for analyzing the empirical residual component for several Ross recovery methods. That component is computed as the difference between the recovered SDF and the Minimum Variance SDF. Column two shows the expected value of the residual component and the p-value for testing the hypothesis that the expected value of the residual component equals zero using a t-test. Column three shows the variance of the residual component.

Analyzing the error component $\hat{\epsilon}_\tau = SDF_\tau - SDF_\tau^{\text{MinVar}}$ **Recovery Method**	$E(\epsilon_\tau)$ (p-value) of testing $E(\epsilon_\tau) = 0$:	$Var(\epsilon_\tau)$
Ross Basic $\pi_{i,j} > 0$	0.308 (0.000)	2.602
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.030 (0.000)	0.016
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.056 (0.000)	0.019
Ross Stable Do not use transition state prices	0.005 (0.450)	0.017
Power Utility with $\gamma = 3$	0.000 (0.919)	0.001

Again, I perform two robustness studies. My findings do not qualitatively change much when reducing the state space by approximately 20%. When removing recession periods, I cannot reject $E(\epsilon_\tau) = 0$ for the Power Utility residual component anymore. It seems that when leaving out recession periods, the expected value of the Power Utility SDF with a risk-aversion parameter of $\gamma = 3$ significantly decreases, which increases its distance to the expected value of the MinVar SDF. Robustness result tables are given in Appendix A.2.

C. Do recovered SDFs explain the transitory SDF component?

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

$$\pi_{i,j} = SDF_{i,j} \cdot p_{i,j} = SDF_{i,j} \cdot \frac{p_{i,j}}{p^{true}_{i,j}} \cdot p^{true}_{i,j} = SDF_{i,j} \cdot SDF^P_{i,j} \cdot p^{true}_{i,j}, \tag{21}$$

with $SDF^P_{i,j}$ being the ratio of recovered and true probabilities. The authors interpret the term $SDF^P_{i,j}$ as a permanent component of the stochastic discount factor, which could be caused by permanent shocks to the macro-economy. They label the Ross recovery stochastic discount factor SDF in contrast as the transitory stochastic discount factor component. They further note that Ross (2015) will only recover the true probabilities if $SDF^P_{i,j} = 1$. As I use this interpretation in my empirical work, I would like to know if the implicit assumption made by Ross (namely, that the permanent component of the stochastic discount factor is one) might contribute to the empirical failure of the recovery theorem. Empirically, Bakshi et al. (2017) suggest that $SDF^P_{i,j}$ exhibits substantial variation and does not equal one. The authors further point out that the transitory stochastic discount factor component (which should equal our usual SDF from Ross recovery) can be found empirically, for date τ, as:

$$SDF^{\mathrm{T}}_{\tau} = \frac{1}{Rf_\tau} \cdot \frac{F_{\tau+,\infty}}{F_{\tau,\infty}}, \tag{22}$$

where Rf_τ is the 30-day risk-free return at date τ, $F_{\tau,\infty}$ is the price of a 30 year Treasury bond future at date τ, and $F_{\tau+,\infty}$ is the price of a 30 year Treasury bond future 30 calendar days after date τ. To uncover if recovered stochastic discount factors empirically equal the transitory component, I regress realized SDFs implied by a particular recovery approach on the empirical transitory component, which I construct as in Equation 22:

$$SDF^{\mathrm{Recovery}}_{\tau} \; = \; \beta_0 + \beta_1 SDF^{\mathrm{T}}_{\tau} + \epsilon_\tau, \quad \tau = 1, ..., \mathcal{T}. \tag{23}$$

Results for that regression are given in Table III. For Ross Basic, I again find a large negative intercept and a large positive slope coefficient. For Ross Bounded, Ross Unimodal, Ross Stable, and Power Utility I find positive intercepts as well as slopes lower than one. The adjusted R^2 is about zero for all these methods, so they all exhibit very low similarity with the transitory component. Also, I can reject the hypothesis of the intercept being equal to zero and the slope being equal to one for all methods. However, that result does not discredit the Power Utility method, as there is no theoretical reason why its stochastic discount factor should equal the transitory component. The story is different for the Ross recovery methods. Their stochastic discount factors should theoretically match the transitory SDF component, yet empirically do not, which is further evidence why the Ross recovery theorem does not work well empirically.

Table III. **Stochastic discount factors and the transitory SDF component.** I present the results from regressing the time series of realized stochastic discount factors for different recovery methods (including Power Utility) on the time series of the empirical transitory component of the stochastic discount factor.

$SDF_\tau = \beta_0 + \beta_1 SDF_\tau^{\mathrm{T}} + \epsilon_\tau$ **Recovery Method**	Intercept β_0	Slope β_1	p-value β_0=0 and β_1=1	adj. R^2
Ross Basic $\pi_{i,j} > 0$	4.30	-3.01	0.000	0.000
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.11	-0.09	0.000	-0.002
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.97	0.08	0.000	-0.002
Ross Stable Do not use transition state prices	0.88	0.12	0.000	0.002
Power Utility with $\gamma = 3$	0.96	0.03	0.000	-0.003

My main findings in this study are robust to (i) reducing the state space by approximately 20% and (ii) removing recession periods. Robustness result tables are given in Appendix A.3.

V. Recovered SDFs and and Risk-Neutrality

Based on graphical evidence, Jackwerth and Menner (2018) argue that the failure of the constraint Ross recovery methods is caused by SDFs that are close to risk-neutrality. I verify this observation by testing if pseudo risk-neutral returns come from the recovered distributions. First, I generate 10,000 realizations of a risk-neutral economy. For each realization, I draw 380 pseudo risk-neutral returns from the risk-neutral cdfs. I plug in those pseudo returns into the recovered cumulative distribution functions to obtain 380 percentile value u_τ, one for each sample day τ. I then test the percentiles u_τ for uniformity using the four introduced density tests. For each of the 10,000 economies, I report whether uniformity of u_τ is rejected at the 5% level. Table IV shows the rejection rates for all recovery methods.

Table IV. **Test for equality to the risk-neutral distribution.** I generate 10,000 realizations of a risk-neutral economy. For each realization, I draw 380 pseudo risk-neutral returns. I report the rejection rates for the hypothesis that the pseudo returns were drawn from another particular recovery method's physical distributions. I use the Berkowitz test, two versions of the Knüppel test (3 and 4 moments), and the KS test, each at a 5% significance level.

H0: $p_\tau = \hat{q}_\tau$ **Recovery Method**	**Berkowitz** rej. rate	**Kolmogorov-Smirnov** rej. rate	**Knüppel 3 moments** rej. rate	**Knüppel 4 moments** rej. rate
Risk-Neutral	5%	5%	2%	5%
Ross Basic $\pi_{i,j} > 0$	94%	91%	88%	91%
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	28%	5%	4%	26%
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	83%	7%	20%	91%
Ross Stable No transition state prices	5%	5%	2%	4%
Power Utility with $\gamma = 3$	70%	65%	51%	58%

A high rejection rate indicates that the recovered cumulative distribution functions do significantly differ from the risk-neutral ones. Unsurprisingly, I find a low 5% rejection rate (2% for the Knüppel 3 moments test) for the hypothesis that risk-neutral returns indeed come from the risk-neutral distribution. For Ross Basic and Power Utility the risk-neutrality hypothesis is rejected for most of the realized economies. Those methods clearly recover physical probabilities that do not equal risk-neutral ones. For Ross Unimodal, I find mixed results. While the Berkowitz test and the Knüppel 4 moments test mostly find a rejection of risk-neutrality, the Kolmogorov-Smirnov test only rejects in 7% of the cases and the Knüppel 3 moments test in only 20% of the cases. For Ross Bounded, rejection rates are all below 28%, which indicates that its recovered physical probabilities are close to risk-neutrality. With rejection rates at 5% or below, Ross Stable essentially recovers risk-neutral probabilities.

VI. Conclusion

In this paper, I investigate the empirical properties of the recovered stochastic discount factors. I implement several methods of Ross recovery starting with 'Ross Basic', which requires minimal assumptions, and continue with the methods 'Ross Bounded', and 'Ross Unimodal', for which I subsequently add plausible economic constraints. An alternative implementation of Ross recovery is 'Ross Stable', which allows recovery without requiring *transition state prices* as an ingredient. I use a Power Utility stochastic discount factor with risk-aversion parameter $\gamma = 3$ as a benchmark method.

For my 380 sample days, I compute a time series of stochastic discount factor realizations, one for each recovery method.

First, I investigate whether recovered realized stochastic discount factors are able to price the S&P 500 and the risk-free asset using the standard pricing equation. All Ross recovery methods fail in doing so. The Power Utility SDF is able to individually (yet not simultaneously) price both assets. I further introduce the Minimum Variance stochastic discount factor, which is constructed as the SDF with the lowest variance that perfectly prices both assets, as a baseline method.

Next, I decompose the recovered stochastic discount factors into the Minimum Variance stochastic discount factor plus a residual SDF component. For any valid stochastic discount factor, that component must be orthogonal to the assets being priced and must equal zero in expectation. For each recovery method, I then test whether that component has an expected value of zero. I can reject that hypothesis for all methods but Ross Stable and Power Utility.

I further compute the variance of the residual SDF component. A low variance implies a high resemblance of the recovered SDF to the Minimum Variance SDF and can reduce the variance for asset pricing. Yet, all Ross recovery methods deliver a much higher variance of that residual component compared to the Power Utility method. Overall, I find that the Power Utility SDF compares much better to the theoretically motivated Minimum Variance SDF.

In my next study, I check on an implicit assumption of Ross, which sets the permanent stochastic discount factor component to one and leads to the recovery of only the transitory stochastic discount factor component. Comparing the recovered transitory component to an empirically estimated transitory component as in Bakshi et al. (2017), I find that the two approaches do not coincide. This provides further evidence that the Ross recovery theorem does not work well empirically.

Finally, I implement a simulation study to show that those Ross recovery methods for which additional plausible economic constraints are added, namely Ross Bounded and Ross Unimodal, recover physical distributions that are close to risk-neutral ones. This also holds for the Ross Stable method. For Ross Basic and the Power Utility method, recovered distributions significantly differ from risk-neutral ones.

My study on the properties of recovered stochastic discount factors provides further evidence that the Ross recovery theorem fails empirically and adds to the understanding of the reason behind this finding. I further contribute to the literature on the decomposition of stochastic discount factors by linking a direct application of a Ross recovery theorem using index options to the research that focuses on Ross recovery from a long-term risk pricing perspective.

REFERENCES

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

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

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

Bakshi, G., Chabi-Yo, F., and Gao, X. (2017). A Recovery that We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem. *Review of Financial Studies*, 31(2):532—555.

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

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

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

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

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

Cuesdeanu, H. and Jackwerth, J. C. (2018). The Pricing Kernel Puzzle in Forward Looking Data. *Review of Derivatives Research, forthcoming.*

Dillschneider, Y. and Maurer, R. (2018). Functional Ross Recovery: Theoretical Results and Empirical Tests. *Working Paper, Goethe University Frankfurt.*

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

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

Jackwerth, J. C. and Menner, M. (2018). Does the Ross Recovery Theorem Work Empirically? *Working Paper, University of Konstanz.*

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

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

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

Qin, L., Linetsky, V., and Nie, Y. (2016). Long Forward Probabilities, Recovery and the Term Structure of Bond Risk Premiums. *The Review of Financial Studies, forthcoming.*

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

Walden, J. (2017). Recovery with Unbounded Diffusion Processes. *Review of Finance*, 21(4):1403–1444.

Yao, X. (2018). What Information Can Ross' Recovery Theory Recover? - Evidence from S&P 500 Index Option Prices. *Working Paper, University of Connecticut.*

Appendix

A.1. Robustness: Pricing the S&P 500 and the risk-free asset

Table V shows results of my study on the pricing ability of SDFs when applying Ross recovery on a state space for which the number of states is reduced by approximately 20%.

Table V. Test if SDFs price the S&P 500 and the risk-free asset when reducing the original state space for Ross recovery by approximately 20%. The first column shows the values of $E\left[SDF \cdot R\right]$ for S&P 500 returns R and the p-values for testing this expectation against unity. The second column shows the values of $E\left[SDF \cdot R_f\right]$ for risk-free returns R_f and the again p-values for testing this expectation against unity. The third column shows the p-values of a Wald test that simultaneously tests both expectation values against unity.

SDF from Method	$E\left[SDF \cdot R\right]$ (p-value - unity)	$E\left[SDF \cdot R_f\right]$ (p-value - unity)	**Joint, Wald Test** (p-value - joint unity)
Ross Basic $\pi_{i,j} > 0$	2.121 (0.203)	2.213 (0.212)	- (0.004)
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.024 (0.000)	1.019 (0.000)	- (0.000)
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.046 (0.000)	1.041 (0.000)	- (0.000)
Ross Stable No transition state prices	1.036 (0.000)	1.032 (0.000)	- (0.000)
Power Utility with $\gamma = 3$	0.999 (0.910)	1.000 (0.970)	- (0.028)
MinVar	1.000 (1.000)	1.000 (1.000)	- (1.000)

Table VI shows the results of my analysis on the pricing ability of SDFs when excluding recession periods from my sample.

Table VI. **Test if SDFs price the S&P 500 and the risk-free asset when excluding recession periods from the original sample.** The first column shows the values of $E\left[SDF \cdot R\right]$ for S&P 500 returns R and the p-values for testing this expectation against unity. The second column shows the values of $E\left[SDF \cdot R_f\right]$ for risk-free returns R_f and the again p-values for testing this expectation against unity. The third column shows the p-values of a Wald test that simultaneously tests both expectation values against unity.

SDF from Method	$E\left[SDF \cdot R\right]$ (p-value - unity)	$E\left[SDF \cdot R_f\right]$ (p-value - unity)	**Joint, Wald Test** (p-value - joint unity)
Ross Basic $\pi_{i,j} > 0$	1.306 (0.001)	1.309 (0.001)	- (0.002)
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.033 (0.000)	1.027 (0.000)	- (0.001)
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.061 (0.000)	1.055 (0.000)	- (0.000)
Ross Stable No transition state prices	1.009 (0.007)	1.003 (0.324)	- (0.007)
Power Utility with $\gamma = 3$	0.994 (0.216)	0.993 (0.292)	- (0.001)
MinVar	1.000 (1.000)	1.000 (1.000)	- (1.000)

A.2. Robustness: The orthogonal SDF component

Table VII shows the results of my analysis on the orthogonal SDF component when applying Ross recovery on a state space for which the number of states is reduced by approximately 20%. The main results do qualitatively not change much. I cannot reject $E(\epsilon_\tau) = 0$ for Ross Basic anymore. However, when looking at the expected value and the variance of ϵ_τ it becomes clear that the Ross Basic SDF is now even more chatic and less reasonable than in the original study. In this robustness study, Ross Stable performs worse than in the original study.

Table VII. **Analyzing the residual SDF component of recovered SDFs when reducing the original state space for Ross recovery by approximately 20%.** I present the results for analyzing the empirical residual component for several Ross recovery methods. That component is computed as the difference between the recovered SDF and the Minimum Variance SDF. Column two shows the expected value of the residual component and the p-value for testing the hypothesis that the expected value of the residual component equals zero using a t-test. Column three shows the variance of the residual component. I reduce the dimension of the original state-space by 20%.

Analyzing the error component $\hat{\epsilon}_\tau = SDF_\tau - SDF_\tau^{\text{MinVar}}$ **Recovery Method**	$E(\epsilon_\tau)$ (p-value) of testing $E(\epsilon_\tau) = 0$:	$Var(\epsilon_\tau)$
Ross Basic $\pi_{i,j} > 0$	1.212 (0.212)	356.9
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.019 (0.002)	0.014
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.041 (0.000)	0.018
Ross Stable Do not use transition state prices	0.032 (0.000)	0.022
Power Utility with $\gamma = 3$	0.000 (0.880)	0.001

Table VIII shows the results of my analysis on the orthogonal SDF component when excluding recession periods from my sample. Results for the Ross Recovery SDFs barely change. The Power Utility SDF performs worse when looking at the expected value of its residual component (an expectation of -0.007 and a p-value of 0.007).

Table VIII. **Analyzing the residual SDF component of recovered SDFs when removing recession periods.** I present the results for analyzing the empirical residual component for several Ross recovery methods. That component is computed as the difference between the recovered SDF and the Minimum Variance SDF. Column two shows the expected value of the residual component and the p-value for testing the hypothesis that the expected value of the residual component equals zero using a t-test. Column three shows the variance of the residual component. I remove recession periods from the original sample.

Analyzing the error component $\hat{\epsilon}_\tau = SDF_\tau - SDF_\tau^{\text{MinVar}}$ **Recovery Method**	$E(\epsilon_\tau)$ (p-value) of testing $E(\epsilon_\tau) = 0$:	$Var(\epsilon_\tau)$
Ross Basic $\pi_{i,j} > 0$	0.308 (0.000)	2.772
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	0.027 (0.003)	0.028
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	0.055 (0.000)	0.030
Ross Stable Do not use transition state prices	0.003 (0.780)	0.031
Power Utility with $\gamma = 3$	-0.007 (0.007)	0.003

A.3. Robustness: Equality to the transitory SDF component

Table IX shows the results of my analysis on the equality of recovered SDFs and the transitory SDF component when applying Ross recovery on a state space for which the number of states is reduced by approximately 20%.

Table IX. **Stochastic discount factors and the transitory SDF component when reducing the original state space for Ross recovery by approximately 20%.** I present the results from regressing the time series of realized stochastic discount factors for different recovery approaches (including Power Utility) on the time series of the empirical transitory component of the stochastic discount factor. I reduce the dimension of the original state-space by 20%.

$SDF_\tau = \beta_0 + \beta_1 SDF_\tau^{\mathrm{T}} + \epsilon_\tau$ **Recovery Method**	Intercept β_0	Slope β_1	p-value β_0=0 and β_1=1	adjusted R^2
Ross Basic $\pi_{i,j} > 0$	26.29	-24.20	0.340	-0.001
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.07	-0.06	0.000	-0.002
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.19	-0.16	0.000	0.000
Ross Stable Do not use transition state prices	0.54	0.49	0.000	0.010
Power Utility with $\gamma = 3$	0.96	0.03	0.000	-0.003

Table X shows the results of my analysis on the equality of recovered SDFs and the transitory SDF component when excluding recession periods from my sample.

Table X. Stochastic discount factors and the transitory SDF component when removing recession periods from the original sample. I present the results from regressing the time series of realized stochastic discount factors for different recovery approaches (including Power Utility) on the time series of the empirical transitory component of the stochastic discount factor. I remove recession periods from the original sample.

$SDF_\tau = \beta_0 + \beta_1 SDF_\tau^{\mathrm{T}} + \epsilon_\tau$ **Recovery Method**	Intercept β_0	Slope β_1	p-value β_0=0 and β_1=1	adjusted R^2
Ross Basic $\pi_{i,j} > 0$	5.34	-4.05	0.000	0.002
Ross Bounded $\pi_{i,j} > 0$, rowsums $\in [0.9, 1]$	1.29	-0.27	0.000	0.009
Ross Unimodal $\pi_{i,j} > 0$ and unimodal, rowsums $\in [0.9, 1]$	1.24	-0.19	0.000	0.000
Ross Stable Do not use transition state prices	0.97	0.03	0.000	-0.002
Power Utility with $\gamma = 3$	1.22	-0.24	0.000	0.000

A.4. The orthogonal component of the SDF

Given that the stochastic discount factor $SDF^\star = SDF^{\text{MinVar}} + \epsilon^\star$ has a component $\epsilon^\star$ that is orthogonal to the return R_A of an asset A and that $E(\epsilon^\star) = 0$ and $Cov\left(SDF^{\text{MinVar}}, \epsilon^\star\right) = 0$, it follows for the variance of pricing asset A with the stochastic discount factor $SDF^\star$:

$$
\begin{aligned}
&Var\left(SDF^\star \cdot R_A\right) = Var\left(\left(SDF^{\text{MinVar}} + \epsilon^\star\right) \cdot R\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot Cov\left(SDF^{\text{MinVar}} \cdot R_A, \epsilon^\star \cdot R_A\right) + Var\left(\epsilon^\star \cdot R_A\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot E\left(SDF^{\text{MinVar}} \epsilon^\star R_A^2\right) - 2 \cdot E\left(SDF^{\text{MinVar}} R_A\right) \underbrace{E\left(\epsilon^\star R_A\right)}_{=0} \\
&\qquad + Var\left(\epsilon^\star \cdot R_A\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot Cov\left(SDF^{\text{MinVar}} \epsilon^\star, R_A^2\right) + 2 \cdot \underbrace{E\left(SDF^{\text{MinVar}} \epsilon^\star\right)}_{=Cov\left(SDF^{\text{MinVar}}, \epsilon^\star\right)=0} E\left(R_A^2\right) \\
&\qquad + Var\left(\epsilon^\star \cdot R_A\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot Cov\left(SDF^{\text{MinVar}} \epsilon^\star, R_A^2\right) + Var\left(\epsilon^\star \cdot R_A\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot Cov\left(SDF^{\text{MinVar}} \epsilon^\star, R_A^2\right) + E\left({\epsilon^\star}^2 \cdot R_A^2\right) - \left(\underbrace{E\left(\epsilon^\star \cdot R_A\right)}_{=0}\right)^2 \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + 2 \cdot Cov\left(SDF^{\text{MinVar}} \epsilon^\star, R_A^2\right) + Cov\left({\epsilon^\star}^2, R_A^2\right) + E\left({\epsilon^\star}^2\right) \cdot E\left(R_A^2\right) \\
&= Var\left(SDF^{\text{MinVar}} \cdot R_A\right) + Cov\left(\epsilon^\star \cdot \left(\epsilon^\star + 2 \cdot SDF^{\text{MinVar}}\right), R_A^2\right) + Var\left(\epsilon^\star\right) \cdot E\left(R_A^2\right)
\end{aligned}
\tag{24}
$$

Setting $R_A = 1$ gives

$$
Var\left(SDF^\star\right) = Var\left(SDF_\tau^{\text{MinVar}} + \epsilon_\tau^\star\right) = Var\left(SDF^{\text{MinVar}}\right) + Var\left(\epsilon^\star\right). \tag{25}
$$

Equation 25 demonstrates that a lower variance of the orthogonal component $\epsilon^\star$ implies a stochastic discount factor $SDF^\star$ that resembles the Minimum Variance SDF. Although the covariance term $Cov\left(\epsilon^\star \cdot \left(\epsilon^\star + 2 \cdot SDF^{\text{MinVar}}\right), R_A^2\right)$ in Equation 24 would have to be studied in more detail, a lower variance of $\epsilon^\star$ reduces the value of $Var\left(\epsilon^\star\right) \cdot E\left(R_A^2\right)$ and indicates a lower variance for pricing the asset A.