# Recovery Theory in a Nonstationary Economy*

Steven Heston
University of Maryland

Kris Jacobs
University of Houston

Fernando Soares
London School of Economics

March 13, 2026

**Abstract**

This paper generalizes the Ross (2015) Recovery Theorem to allow for a persistent non-stationary component in asset prices. The stochastic discount factor is given by the reciprocal gross return on a long-term security that pays a power of the market index. This setup allows for growing economies with nonstationary state variables. By imposing path-independence, we can implement this theory and extract the pricing kernel from long-term power securities. Our empirical application synthesizes the power security using a novel dataset of long-term (10-year) options on the S&P 500 index. Our estimate of the stochastic discount factor has plausible time-series properties and is able to explain the equity premium, the long-bond premium, and a negative variance premium. The decomposition of the index equity premium across the return state space is mostly consistent with the data for our estimate of the stochastic discount factor, unlike existing equilibrium models (Beason and Schreindorfer, 2020).

*Heston: sheston@umd.edu; Jacobs: kjacobs@bauer.uh.edu. We thank Jon Danielsson and Jean-Pierre Zigrand for resources from the Systemic Risk Centre at the London School of Economics. We are also grateful for comments from Jaroslav Borovicka, John Crosby, Juhani Linnainmaa, Karamfil Todorov, and participants at the Conference on Derivatives and Volatility, the Virtual Derivatives Workshop, and the NYU Finance and Risk Engineering seminar.

## 1 Introduction

Option theory has long aspired to use market prices to infer information about observable (historical) probabilities. Following the insights of Breeden and Litzenberger (1978), the extraction of the risk-neutral return distribution from equity option prices has become an essential part of the financial economics toolkit, and this approach has been used in a variety of empirical applications. However, in a seminal contribution, Ross (2015) forcefully argued that we made limited progress in reliably identifying the historical distribution from the extracted risk-neutral one, or equivalently decomposing the risk-neutral probabilities in historical probabilities and a pricing kernel. The dominant approach in financial economics to implement this decomposition relies on a parametric economic model, and Ross rightly pointed out that there is no consensus yet as to what this model should be.

Ross (2015) instead proposes an an ambitious extension of this research program to "recover" the entire distribution of returns on the spot asset. His approach is elegant because it makes a simple, parsimonious, model-free assumption that the pricing kernel depends on the state, but not the path of the state. This assumption nests time-separable expected utility, and also allows for more general Kreps and Porteus (1978) recursive preferences. However, several authors have rightly pointed out that Ross' assumptions are too stringent, or equivalently that they rule out many or most realistic models (Borovička, Hansen, and Scheinkman, 2016; Jensen, Lando, and Pedersen, 2019)). Most importantly, this assumption imply a stationary economy and price level. This raises concerns about implementation, because all major U.S. equity indices appear to be nonstationary.

This paper builds on the path-independence assumption in Ross (2015), but our framework allows for a nonstationary component in prices. Like Ross (2015), our approach is motivated by the principle of parsimony and does not rely on a parametric model. Because the empirical evidence shows that the pricing kernel has a large permanent nonstationary component Alvarez and Jermann (2005), we allow for a nonstationary spot price, which implies that the log-pricing kernel is cointegrated with the spot price. This assumption nests the ICAPM of Merton (1973a) and the production models in Cox, Ingersoll, and Ross (1985) which have nonstationary wealth accumulation in addition to a stationary investment opportunity set. A path-independent pricing kernel is also consistent with power utility (Rubinstein, 1976; Brennan, 1979) and the recursive preferences of (Epstein and Zin, 1989, 1991). Our approach is also very general in that it allows for the presence of other (stationary) state variables, such as (stochastic) volatility and more structural economic variables.

The economy in Ross (2015) can be thought of as a special case of ours, when the power parameter is equal to zero. Another (equivalent) way to characterize the relation between our approach and the existing literature on recovery, and Ross (2015) in particular, is that recovery theory can be thought of as choosing a numeraire. Ross (2015) effectively assumes that prices are stationary in a dollar numeraire. This is highly counterfactual. It also implies that the long-term bond is a growth-optimal investment. In other words, the long-term bond is the numeraire portfolio of Long (1990), and risk premia are determined by the covariation between returns and interest rates. Because interest rates are not volatile, stationary recovery theory implies small risk premia. We use a change of numeraire to identify a different numeraire portfolio. We show that a long-term power security that pays proportional to the gross market return raised to a power is growth-optimal in our nonstationary recovery framework. This long-term power security can be thought of as providing exposure to equity risk, long-term interest rate risk, and volatility risk.

Is the path-independence assumption appropriate and/or realistic? In the end this is an empirical question. Path independence is motivated by the principle of parsimony, that the pricing kernel depends on the same Markovian variables as asset prices. We therefore think of the path-independent pricing kernel as first among equals and an obvious candidate for empirical work. It provides a pragmatic and model-free way to study the properties of the pricing kernel.

We start our analysis with the most general characterization of our approach, which does not require parametric assumptions. We provide a summary of the existing literature on recovery theory, and explain how our approach differs. To provide more intuition for our approach, we then derive our result in two important parametric setups: the Black and Scholes (1973) model and the Heston (1993) square root stochastic volatility model. We provide more intuition for our general result by relating it to the specific outcomes in these parametric examples.

Our setup can be implemented with the help of data on long-term options, and we provide a detailed empirical analysis. We use the approach of Bakshi and Madan (2000) and Bakshi and Kapadia (2003) to replicate the power security using a new dataset that provides a long time series of prices on long-term options for a wide range of moneyness. We provide the results of this replication for different values of the power parameter. We study the resulting time series of power security returns through their descriptive statistics and through their relation with the returns on the forward, bond and variance securities. Intuitively the long-term power security can be replicated through exposure to equity (stock market index) risk, long-term interest rate risk, and (index) volatility risk. The return on the stock market index and the bond return correspond to the return on the power security in two special (nested) cases, which correspond to fixed values of the power parameter. The return on volatility risk captures the nonlinear (concave) dependence of the power security return on the index. We confirm empirically that the returns on the power security are spanned by the returns on these three securities, and that the loadings on the three security returns confirm the theoretical properties of the power security.

Encouraged by these properties of the power security return, we then turn to an analysis of the implied estimates of the pricing kernel. The values of the estimated pricing kernels as a function of index returns are intuitively plausible, and the time series of the pricing kernels are reasonable as a function of economic conditions. Next we consider the framework of Beason and Schreindorfer (2022), who demonstrate that existing consumption-based models attribute a large part of the index equity premium to the extreme left tail of returns. This contrasts with the empirical distribution, where returns below -30% matter very little for the index equity premium. We find that in contrast to consumption-based models, our estimated pricing kernels are much more closely aligned with the data.

Our contribution straddles several literatures that take different approaches to constructing and estimating candidate stochastic discount factors. Besides the literature on recovery theory mentioned above, the most important component of this literature uses equilibrium consumption-or production based representative-agent models. This literature is too extensive to cite in full here. A more closely related approach combines the insights of Breeden and Litzenberger (1978) together with an assumption on the physical return dynamic to estimate conditional and unconditional pricing kernels.[1] Chernov (2003) reverse engineers the pricing kernel based on options on various securities. Ghosh, Julliard, and Taylor (2017) nonparametrically estimate the stochastic discount factor using a relative entropy minimization approach. Finally, Alvarez and Jermann (2005) emphasize the importance of the permanent component of the pricing kernel and study its properties.

The paper proceeds as follows. We present the general result in Section 2. Section 3 discusses two important parametric examples. In Section 4, we discuss the data sources, the construction of the synthetic assets, and the properties of the security returns used in the empirical analysis. Section 5 presents and analyzes the estimates of the pricing kernels. Section 6 concludes.

# 2 Recovery Theory in Economies with Growth

We first provide a brief summary of recovery theory and discuss the existing literature. We then present the most general version of our result.

## 2.1 Ross Recovery Theory

Recovery Theory was proposed in Ross (2015). The theory was subsequently extended in Hansen and Scheinkman (2009), Qin and Linetsky (2016), Borovička, Hansen, and Scheinkman (2016), and Jensen, Lando, and Pedersen (2019). Martin and Ross (2019) use recovery theory to study the properties of the long end of the term structure.

[1] See for instance Aït-Sahalia and Lo (1998), Aït-Sahalia and Lo (2000), Jackwerth and Rubinstein (1996), Jackwerth (2000), Rosenberg and Engle (2002). Christoffersen, Heston, and Jacobs (2013), Song and Xiu (2016), Cuesdeanu and Jackwerth (2018), Linn, Shive, and Shumway (2018), and Barone-Adesi, Fusari, Mira, and Sala (2020).

## 2.2 Recovery Theory in Economies with Growth: The General Case

We consider the valuation at time $t$ of an asset $U_t$ that pays out at future times $t+\Delta$ and is valued by a stochastic discount factor $\frac{M_{t+\Delta}}{M_t}$. That is, the value of any asset $U_t$ satisfies $U_t = E_t[U_{t+\Delta}\frac{M_{t+\Delta}}{M_t}]$. In our analytical specification, we often refer to $M_t$ rather than the stochastic discount factor $\frac{M_{t+\Delta}}{M_t}$. To avoid confusion with terminology, we follow Alvarez and Jermann (2005) and refer to $M_t$ as the pricing kernel.

We deliberately do not assume a fully parametric model and make as few assumptions as possible. Instead, we consider an extension of Ross (2015) and assume that the pricing kernel at time $t$, $M_t$, is specified by:

$$M_t = M_0 e^{\beta t} S_t^{\gamma} h(X_t) \tag{1}$$

where $S_t$ is the level of the stock market index and $h(X_t)$ denotes a (possibly nonlinear) function $h(.)$ of a vector of stationary state variables $X_t$. Under the customary assumption of a nonstationary $S_t$, this specification of $M_t$ assumes that it is nonstationary and that its logarithm is cointegrated with the logarithm of the stock market index. The $S_t^{\gamma}$ component in equation (1) corresponds to the persistent component in Alvarez and Jermann (2005). In addition to allowing $S_t$ to be nonstationary, we assume that that for any future $T$, $\frac{S_T}{S_t}$ is independent of $S_t$, i.e., that future index returns do not depend on the index level.

Now consider a security that pays $(S_T)^{\phi}$ at time $T$. We will refer to this security as a power security. In general we will denote its value at time $t$ with a payout at $T$ as $PS_{t,T}(\phi)$ to denote that the value depends on the power parameter $\phi$, but we often drop this notation when presenting computations to simplify notation. At time $t$, the value $PS_{t,T}$ of this power security is given by:

$$PS_{t,T} = E_t[(S_T)^{\phi}\frac{M_T}{M_t}] = e^{\beta(T-t)}E_t[(S_t)^{\phi}\frac{(S_T)^{\phi+\gamma}}{(S_t)^{\phi+\gamma}}\frac{h(X_T)}{h(X_t)}] \tag{2}$$

For $\phi = -\gamma$, the value of the power security therefore reduces to $PS_{t,T} = e^{\beta(T-t)}(S_t)^{\phi}E_t[\frac{h(X_T)}{h(X_t)}]$. Now let $T \to \infty$ in equation (2), and consider the ratio of $PS_{t+\Delta,T\to\infty}$ and $PS_{t,T\to\infty}$. This gives us the following return on an *asymptotic* power security:

$$\frac{PS_{t+\Delta,T\to\infty}}{PS_{t,T\to\infty}} = e^{\beta\Delta}(\frac{S_{t+\Delta}}{S_t})^{\phi}\frac{h(X_t)}{h(X_{t+\Delta})} = e^{\beta\Delta}(\frac{S_t}{S_{t+\Delta}})^{\gamma}\frac{h(X_t)}{h(X_{t+\Delta})} \quad (3)$$

Equation (3) shows that for $\phi = -\gamma$, the gross return on this asymptotic power security is equal to $\frac{M_t}{M_{t+\Delta}}$, the reciprocal of the stochastic discount factor. This result is very general and valid in models with multiple factors, including the volatility and jump factors typically used in reduced-from models, as well as structural and economic factors. The only requirement is that the factors are stationary.

The pricing kernel in equation (1) can be motivated in two ways. From the perspective of recovery theory, the kernel reduces to the assumptions in Ross (2015) when $\gamma = 0$, i.e., in a stationary economy. We discuss this in more detail in Section 3.3 below. From the perspective of more structural consumption-based asset pricing, equation (1) is consistent with a variety of models that imply nonstationary economies such as the ICAPM in Merton (1973a) and Epstein and Zin (1989) preferences.

In summary, we have established that the gross return on an asymptotically long power security is the reciprocal of the stochastic discount factor. This power security can therefore be thought of as Long (1990) log-efficient "numeraire portfolio". All asset prices should be martingales under this numeraire. It is critical that this result does not require parametric assumptions on $S(t)$. However, it is instructive to examine this general result for specific parametric assumptions on $S(t)$, because this relates our result to the existing literature and provides valuable insights into the implications for risk premiums. We now turn to this task.

# 3 Recovery Theory in Economies with Growth: Parametric Examples

We provide more intuition for our result by illustrating how it obtains for two very important parametric specifications, the Black-Scholes Merton model (Black and Scholes, 1973; Merton, 1973b) and the square root stochastic volatility model in Heston (1993). We conclude with a discussion and show how our framework presents an economically meaningful generalization of the setup in Ross (2015).

### 3.1 The Black-Scholes-Merton (BSM) Model

We apply the result in Section 2.3 to the BSM model. The BSM model is of course a logical starting point for any parametric analysis, but there is some potential for confusion with our application, because the i.i.d. assumption in the Merton model drastically simplifies the framework. It is nevertheless instructive to outline the properties of the power security in the context of the geometric random walk assumption in the BSM model.

The risk-neutral and physical dynamics of the spot index $S_t$ are respectively given by:

$$\begin{aligned} dS_t/S_t &= rdt + \sigma dz_t^*, & (4) \\ dS_t/S_t &= \mu dt + \sigma dz_t, & (5) \end{aligned}$$

where $dz^*$ and $dz$ are Wiener process under the risk-neutral physical measures respectively. The Appendix shows that the implied pricing kernel is defined by:

$$M_t = e^{\beta t} S_t^{\gamma} h(X_t) \tag{6}$$

where $\gamma = (\mu - r)/\sigma$ and $\beta = -(\gamma + 1)\mu - \gamma(\gamma + 1)\sigma^2/2$. The value of the power security can be obtained from the Black-Scholes-Merton PDE.

$$PS_{t,T} = S_t^{\phi} e^{d(T-t)} \tag{7}$$

where $d = \sigma^2\phi(\phi - 1)/2 + r(\phi - 1)$. Letting $T \to \infty$ and taking the ratio of $PS_{t+\Delta,T}$ and $PS_{t,T}$, the return on the asymptotic power security in the BSM model is therefore given by:

$$\frac{PS_{t+\Delta,T\to\infty}}{PS_{t,T\to\infty}} = (\frac{S_t}{S_{t+\Delta}})^{\phi} e^{d(T-t)} \tag{8}$$

Moreover, for $\phi = -\gamma$ we get $d = \beta$, and therefore we confirm the general result that $\frac{PS_{t+\Delta,T\to\infty}}{PS_{t,T\to\infty}} = \frac{M_t}{M_{t+\Delta}}$, the reciprocal of the stochastic discount factor.

An important difference between the parametric BSM model studied in this section and the general case in Section 2.3 is that we can obtain the central result under much weaker assumptions, using a power security with any finite maturity $T$. This can be readily seen by starting from equation (7) and computing the ratio for a finite $T$. The intuition for this result can be easily seen by comparing equations (1) and (6). The BSM kernel does not contain any additional stationary variables $X_t$, and therefore it is not necessary to use an asymptotic power security to difference out the impact of these variables. In other words, this result directly follows from the (restrictive) assumption of a geometric random walk for $S_t$ in the BSM model. Any deviation from this assumption will require the use of the asymptotic power security to obtain the main result that the return on the power security is the reciprocal of the stochastic discount factor. We now turn to one such example, which also takes an important place in the option pricing literature.

### 3.2 The Square Root Stochastic Volatility Model

Our next parametric example is the Heston (1993) continuous-time stochastic square root volatility model, which specifies the risk-neutral dynamics of the spot index $S_t$ and its stochastic variance $v_t$ as follows:

$$\begin{aligned} dS_t/S_t &= rdt + \sqrt{v_t}dz_t^{1*}, \\ dv_t &= \kappa^*(\theta^* - v_t)dt + \sigma\sqrt{v_t}dz_t^{2*}, \end{aligned} \tag{9}$$

where $dz^{1*}$ and $dz^{2*}$ are Wiener processes with correlation coefficient $\rho$. We again consider the simplest case with constant risk-free rate $r$. The physical dynamic has the same functional form as the risk-neutral dynamic:

$$
\begin{aligned}
dS_t/S_t &= [r + \mu v_t]\, dt + \sqrt{v_t} dz_t^1, \\
dv_t &= \kappa(\theta - v_t)dt + \sigma\sqrt{v_t} dz_t^2,
\end{aligned}
\tag{10}
$$

where $\mu v_t$ is the equity premium as a function of $v_t$, and $dz_1$ and $dz_2$ are Wiener processes under the physical measure. While $\eta$, the variance of variance parameter, and $\rho$, the correlation between $z_1$ and $z_2$, are assumed to be identical to the corresponding parameters in the risk-neutral dynamics, the long-run physical variance $\theta$ and mean reversion $\kappa$ differ from the long-run risk-neutral variance $\theta^*$ and mean reversion $\kappa^*$.

The Appendix shows that the pricing kernel consistent with both risk-neutral (9) and physical dynamics (10) takes the form

$$
M_t = M_0 S_t^{\gamma} \exp\left(\beta t + \eta \int_0^t v_s ds + \xi v_t\right), \tag{11}
$$

where

$$
\begin{aligned}
\beta &= -(1+\gamma)\, r - \xi \kappa^* \theta^*, && (12) \\
\eta &= -\frac{1}{2}\left(\gamma^2 + 2\rho\sigma\gamma\xi + \sigma^2\xi^2\right) - (\mu - \frac{1}{2})\gamma + \kappa\xi, && (13) \\
\gamma &= -\mu - \rho\sigma\xi, && (14) \\
\xi &= \frac{\rho\sigma\mu - \kappa^* + \kappa}{(1-\rho^2)\sigma^2}. && (15)
\end{aligned}
$$

The path-independence condition associated with recovery theory imposes $\eta = 0$ in equation (11), which gives:[2]

$$
M_t = M_0 S_t^{\gamma} \exp\left(\beta t + \xi v_t\right), \tag{16}
$$

[2]The Appendix discusses how the functional form of the $\xi v(t)$ term in equations (11) and (16) is related to the results in Hansen and Scheinkman (2009) and Qin and Linetsky (2016).

We again consider a power security that pays $(S_T)^\phi$ at time $T$. Its value at time $t$ is given by $PS_{t,T} = (S_t)^\phi e^{-r(T-t)+C(T-t;\phi)+D(T-t;\phi)v_t}$, where $C(.)$ and $D(.)$ are the solutions for the generating functions as in Heston (1993). Denoting $\tau = T - t$, we have:

$$\begin{aligned}
C(\tau;\phi) &= r\phi\tau + \frac{\kappa^*\theta^*}{\sigma^2}((\kappa^* - \rho\sigma\phi + d)\tau - 2\log(\frac{1 - g\exp(d\tau)}{1-g})), \\
D(\tau;\phi) &= \frac{\kappa^* - \rho\sigma\phi + d}{\sigma^2}(\frac{1 - g\exp(d\tau)}{1-g}), \\
d &= \sqrt{(\kappa^* - \rho\sigma\phi)^2 + \sigma^2(\phi - \phi^2)}, \\
& \quad g = \frac{\kappa^* - \rho\sigma\phi + d}{\kappa^* - \rho\sigma\phi - d}.
\end{aligned}$$

As $\tau \to \infty$, we have $C(\tau; -\gamma) \to \delta\tau$ and $D(\infty; -\gamma) \to -\xi$. Therefore

$$\frac{PS_{t+\Delta,T\to\infty}}{PS_{t,T\to\infty}} = \frac{(S_{t+\Delta})^\phi e^{-r(T-t-\Delta)+C(T-t-\Delta;\phi)+D(T-t-\Delta;\phi)v_{t+\Delta}}}{(S_t)^\phi e^{-r(T-t)+C(T-t;\phi)+D(T-t;\phi)v_t}} \tag{17}$$

Using equation (16), this confirms that for $\phi = -\gamma$, the gross return on the asymptotic power security is once again $\frac{M_t}{M_{t+\Delta}}$, the reciprocal of the stochastic discount factor.

### 3.3 Discussion

The fundamental theorem of asset pricing asserts the existence of a stochastic discount factor that prices all assets in the absence of arbitrage. The recovery framework in Ross (2015) critically relies on restricting this class of pricing kernels by imposing *transition-independence* (p. 619, equation (10)). Using the notation from equation (1), the transition-independence assumption amounts to $M_t = M_0 e^{\beta t} h(X_t)$. Given this definition, consider the ratio of the stochastic discount factor for the valuation (at time $t$) across two future (time $t + \Delta$) states where $X_{t+\Delta} = A$ and $X_{t+\Delta} = B$:

$$\frac{\frac{M^A_{t+\Delta}}{M_t}}{\frac{M^B_{t+\Delta}}{M_t}} = e^{\beta\Delta}\frac{\frac{h(A)}{h(X_t)}}{\frac{h(B)}{h(X_t)}} = e^{\beta\Delta}\frac{h(A)}{h(B)}. \tag{18}$$

This ratio is independent of the state at time $t$. Transition-independence thus effectively eliminates dependence of valuation on the historical path of $X(t)$ as well as $S(t)$. It enforces the economic condition that preferences depend on "where we are", but not "how we got there" via jumps, diffusions, or any other history.

Note that our proposed kernel in equation (1) is also independent of the path of historical states. It generalizes the specification in Ross (2015) by allowing the pricing kernel to depend on a nonstationary price level, and reduces to the one in Ross (2015) when $\gamma = 0$. Following Heston (2004), we refer to our kernel in equation (1) as the *path-independent* kernel.

The importance of allowing for nonzero $\gamma$ is illustrated by our two parametric examples, the Black-Scholes Merton model (Black and Scholes, 1973; Merton, 1973b) and the square root stochastic volatility model in Heston (1993). Both economies are non-stationary, and this will generally also be the case for other parametric examples such as consumption-based models, which can typically be nested in equation (1). For instance, our proposed kernel in equation (1) is consistent with Epstein and Zin (1989, 1991) and Kreps and Porteus (1978) recursive preferences. Other asset pricing models such as the ICAPM (Merton, 1973a) are also nonstationary. Note also that while our proposed pricing kernel implies that future returns are independent of stock price *levels*, they are not necessarily independent of past returns. The kernel is therefore consistent with models that specify some predictability of returns by lagged returns or other predictors.

From an empirical perspective, most major equity indices display exponential long-term growth. Our nonstationary specification (1) is consistent with this stylized fact and allows the log-pricing-kernel to be cointegrated with the log-stock index. This "permanent component" of the pricing kernel is also supported empirically by Alvarez and Jermann (2005), who show that the the permanent component of the pricing kernel is large and volatile relative to the transitory component $h(X)$.

Finally, to appreciate the importance of the transition-independence and path-independence specifications, consider the kernel for the stochastic volatility model in equation (11). This kernel is path-dependent because of the presence of the integrated history of the variance $v_t$. The kernel can be made path-independent by imposing the restriction $\eta = 0$.

# 4 Data Sources and Data Construction

We present and discuss the data used in the empirical analysis. We rely on a comprehensive dataset on derivative prices. These data are relatively new to the literature and we therefore describe them in some detail. We then explain how to compute the prices of the synthetic assets as well as monthly returns. We provide descriptive statistics for these returns and analyze the returns on the power security.

## 4.1 Data Sources

Our framework requires data on deep out-of-the-money index options with long maturities. The OTCDD Derivatives Data provided by S&P Global are very well suited for this purpose. These data combine exchange data with daily price submissions by major market-making dealers. For options with insufficient exchange data, the OTCDD data aggregate the dealers' individual submissions into consensus prices. This approach yields a much more extensive cross-section of maturities and moneyness than more traditional sources of derivative data.

We rely on the OTCDD data for European calls and puts issued on the S&P 500, which we regard as representative of the market index. These options have moneyness ranging from 50% to 200% and maturities of up to 10 years. Prices are reported at the daily frequency for the 2007-2023 sample period. The call and put S&P 500 options are quoted either OTM or ATM. For each available maturity, there is an ATM call and an ATM put, which can be used to compute the forward price using put-call parity.

To ensure consistency with the option prices, we rely on market index returns computed from the daily closing price for the S&P 500 provided by OTCDD. We compute realized variances as the squared log returns of S&P 500 evaluated at a daily frequency. The top panel in Figure 1 plots the S&P 500 return (left axis) and realized variance (right axis) for non-overlapping monthly intervals. The bottom panel plots the cumulative return on the S&P 500.

To compute risk-free returns, we use the term structure of U.S Treasury yields parameterized according to the methodology of Gürkaynak, Sack, and Wright (2007) and fitted daily.[3] Relying on a parameterized term structure enables us to compute yields with the specific maturities of the options available each day. These yields are then used to compute risk-free returns.

## 4.2 Constructing Synthetic Assets

Our results rely on the prices of two assets that cannot be directly observed, the power security and the long-term variance swap on the market index. We now discuss how to construct these synthetic assets and extract their prices from the options data. Let $C(K;t,T)$ and $P(K;t,T)$ denote the prices at time $t$ of the call and put options issued on the market index $S_t$, with expiration date $T$, and spanning a continuum $K \in \mathbb{R}^+$ of strike prices. In the absence of arbitrage, the fair strike price $V_{t,T}$ of a variance swap issued at $t$ with expiration at $T$ is given by

$$V_{t,T} \equiv E_t^* \left[ \int_t^T \left( \frac{dS_t}{S_t} \right)^2 \right] = 2R_{t,T}^f \left( \int_0^{F_{t,T}} \frac{P(K;t,T)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{C(K;t,T)}{K^2} dK \right) \tag{19}$$

where the $*$ superscript denotes expectations under the risk-neutral measure. The price of the power security is then equal to

$$PS_{t,T}(\phi) \equiv \frac{1}{R_{t,T}^f} E_t^* \left[ S_T^{\phi} \right] = \frac{\left(E_t^* \left[S_T\right]\right)^{\phi}}{R_{t,T}^f} + \phi\left(\phi - 1\right) \left( \int_0^{F_{t,T}} \frac{P(K;t,T)}{K^{2-\phi}} dK + \int_{F_{t,T}}^{\infty} \frac{C(K;t,T)}{K^{2-\phi}} dK \right) \tag{20}$$

where $F_{t,T}$ is the forward price and $R_{t,T}^f$ denotes the maturity specific risk-free rate.

Our implementation discretizes these integrals over the available strike prices and approximates them numerically. The OTCDD dataset allows for small discretization intervals since it provides 15 strikes for the S&P 500 options. We use Simpson's rule to further enhance numerical accuracy.

[3] These data are available at https://www.federalreserve.gov/data/nominal-yield-curve.htm.

### 4.3 Computing Monthly Returns

Our empirical analysis uses monthly returns. Let $t$ and $t+1$ be two dates one month apart. The return on the power security is defined as

$$R^P_{t+1,\tau}(\phi) = \frac{PS_{t+1,t+\tau}(\phi)}{PS_{t,t+\tau}(\phi)}$$

Similarly, we define the gain on the variance swap as:

$$R^{VAR}_{t+1,\tau} = V_{t+1,t+\tau} + RV_{t,t+1} - V_{t,t+\tau}$$

where $\tau = T - t$ is the maturity expressed in months and $RV_{t,t+1}$ is the realized variance over a monthly interval. The forward price is the price of a security that delivers the level of the market index at expiration. Under the assumption of no-arbitrage, its returns is given by

$$R^{FWD}_{t+1,\tau} = \frac{R^f_{t,t+\tau} F_{t+1,t+\tau}}{R^f_{t+1,t+\tau} F_{t,t+\tau}}$$

It should be noted that the forward return is reduced to the spot market index return $R^{SPX}_{t+1}$ in the absence of dividends. Finally, we use the interpolated risk-free rates to calculate the return on bonds:

$$R^B_{t+1,\tau} = \frac{R^f_{t,t+\tau}}{R^f_{t+1,t+\tau}}$$

As with other sources of derivative data, the OTCDD options have fixed expiration dates, which implies that their maturities get progressively shorter over time until a rollover date is reached. For long-term options, the rollover occurs annually on the second week of December. Rollover frequency is higher for shorter maturities. Since the returns need to match the available option data, their maturities also get progressively shorter between rollover dates. However, because the OTCDD maturities are typically much larger than the interval between their respective rollover dates, the returns can be treated as having a constant maturity. For expositional convenience, we refer to the maturity of a return series as the value of $\tau$ in the month subsequent to rollover. For instance, when a return series is said have to have a maturity of 120 months, it has $\tau = 120$ in January and $\tau = 109$ in December.

### 4.4 Descriptive Statistics

Table 1 reports descriptive statistics for the monthly forward returns, bond returns, and variance swap gains. We select 10 years as our baseline long maturity, but we also report on the 6-month, 3-year and 9-year maturity, as well as the correlations between maturities. Panel A of Figure 2 plots the time series of the three securities for the 10-year maturity. As expected, the average forward return exceeds the bond return irrespective of maturity. Because the variance process typically has a negative correlation with the market index, the variance swap gain is on average negative. The returns on the forward and the variance swap are considerably more volatile than the bond return.

Consistent with existing evidence, Table 1 shows that the forward returns display negative skewness at all maturities. Conversely, the bond returns and the variance swap gains are mostly positively skewed. While all return distributions display excess kurtosis, kurtosis is significantly higher for the variance swap gains. Panel A of Figure 2 visually confirms these patterns: the outliers in the time series of the variance swap gains are both larger and more frequent. The correlation matrices in Table 1 show that the 9- and 10-year maturities are highly correlated for all three securities, suggesting that the 10-year maturity is a good approximation to an asymptotically long maturity.

Table 2 reports on the distribution of the time series of the power security return. We report results for different values of the $\phi$ parameter and different maturities. Table 2 shows that conditional on a value of $\phi$, the return distribution is similar across maturities. The correlations in Panel A of Table 3, which are for $\phi = 0.5$, confirm this. For all positive $\phi$ values, the power security return displays negative skewness. Excess kurtosis is moderately positive and peaks at intermediate values of $\phi$. Panel B of Figure 2 plots the time series of the power security return for selected values of $\phi$ for our baseline 10-year maturity. As expected, the time series is more variable as $\phi$ increases, which can also be gleaned from Table 2.

Panel B of Table 3 reports the correlations between the power security and the other three securities for the baseline 10-year maturity. We again report on power security returns for different values of $\phi$. Not surprisingly, the correlation between the power security and the forward returns is increasing in $\phi$ and is positive for positive $\phi$. On the other hand, the correlation between the power security and the bond returns is decreasing in $\phi$ and it switches from positive to negative for sufficiently high $\phi$. The correlation between the power security return and the variance swap gain is negative in the interval $0 < \phi < 1$ and it varies non-monotonically with $\phi$, reaching its minimum around $\phi = 0.5$. The forward is negatively correlated with the bond and the variance swap, while the latter two are positively correlated.

### 4.5 Interpreting the Power Security Return

To provide more insight into the properties of the power security, we analyse the risk factors driving its returns. We identify the individual contributions of the forward log return, the bond log return, and the variance swap gain to the power security log return using the following regression:

$$\ln R_t^P(\phi) = \alpha + \beta^{FWD} \ln R_t^{FWD} + \beta^B \ln R_t^B + \beta^{VAR} R_t^{VAR} + \epsilon_t$$

Table 4 reports the regression results for a maturity of 10 years and selected $\phi$ values. As a general result, the regression has an adjusted $\mathrm{R}^2$ nearly indistinguishable from one regardless of the choice of $\phi$, providing an intuitive interpretation of the power security as a portfolio of the three securities considered with portfolio weights given by the regression coefficient estimates. Table 4 indicates that the weights depend on the value of $\phi$.

Figure 3 provides additional insight into the portfolio interpretation of the power security. We plot the portfolio weights over a fine grid of $\phi$ values. As expected, the weight on the forward increases with the value of $\phi$, while the weight on the bond decreases. Two special cases are worth highlighting: the portfolio consisting solely of one unit of the forward when $\phi = 1$, and the portfolio consisting solely of one unit of the bond when $\phi = 0$. For the interval $0 < \phi < 1$, the portfolio consists of a combination of all three assets, with a negative weight on the variance swap and positive weights on the forward and bond. When $\phi = 0.5$, the variance swap weight reaches its most negative value and the other two assets have roughly equal weights.

# 5 The Stochastic Discount Factor and Risk Premiums

We present estimates of the stochastic discount factor based on the power security returns. We discuss the properties of the time series of the recovered stochastic discount factor, and investigate how well the stochastic discount factor performs in explaining the market equity premium.

## 5.1 Estimates of the Stochastic Discount Factor

We estimate the SDF based on our computations for the power security returns. The analytical results in Section 2 posit that for an arbitrary horizon $\Delta$, the SDF $\frac{M_{t+\Delta}}{M_t}$ can be computed as the reciprocal of the gross power security return $\frac{PS_{t+\Delta,T\to\infty}}{PS_{t,T\to\infty}}$. As mentioned in Section 4, we believe that the 10-year maturity is sufficiently long to approximate the asymptotic behavior of power security returns, rendering our option data adequate for the purpose of estimating the SDF. Consistent with our implementation in Section 4.3, we take $\Delta$ as a monthly interval and estimate the monthly SDF time series $M_{t,t+1}(\phi)$ as the reciprocal of the 10-year power security return $R^P_{t+1}(\phi)$. Because the power security return is a function of $\phi$, this also applies to the SDF time series.

Figure 4 plots the time series of the SDF estimates for different values of $\phi$ and Table 5 reports descriptive statistics for these time series. The estimated fluctuations in the marginal utility of the representative agent seem plausible. More formally, to assess the ability of the estimated SDF to explain returns on a given set of assets, we follow Hansen and Jagannathan (1991) and verify whether the SDF is sufficiently variable. Hansen and Jagannathan (1991) show that the ratio of the standard deviation and the mean of the SDF exceeds the Sharpe ratio of the candidate assets. For $\phi = 0.75$, this means that the estimated SDF can explain assets with an annual Sharpe ratio of 0.145. This is a multiple of the Sharpe ratio implied by a simple consumption-based model with power utility.

## 5.2 Decomposing the Equity Premium

Beason and Schreindorfer (2022) illustrate the shortcomings of prominent consumption-based models by documenting the source of the equity index premium across the return state space in these models and comparing it to the data. They find that while in the empirical distribution, returns below -30% matter very little for the index equity premium, existing consumption-based models attribute a large part of the premium to the extreme left tail. We show that our estimates of the SDF based on the return on the power security perform much better along this dimension.

### 5.2.1 Empirical Physical Density and Recovered SDF

To study the equity premium, we require estimates of the physical and risk-neutral densities of the market index return. As in Beason and Schreindorfer (2022), we compute the physical density $f(R)$ based on the relative frequency of observations within each interval of an evenly-spaced grid of returns. We also follow Beason and Schreindorfer (2022) in calculating the market index return $R_t^{SPX}$ on overlapping one-month intervals (28 trading days) sampled at a daily frequency to ensure a large number of observations and a more efficient estimation of the density function despite our relatively short sample. Denote by $N$ be number of return observations and by $R^k$ the $k^{th}$ node of a grid where the nodes are separated by a distance $2\Delta$:

$$f(R^k) = \frac{1}{N}\sum_{t=1}^{N} \mathbf{1}(R^k - \Delta < R_t^{SPX} \leq R^k + \Delta)$$

We verified that results are almost identical when using a kernel density estimator. In a similar fashion, we also use a non-parametric approach for computing the recovered SDF as a function of the market index return:

$$M^{MOD}(R^k;\phi) = \frac{\sum_{t=1}^{N} M_{t,t+1}^{daily}(\phi)\mathbf{1}(R^k - \Delta < R_t^{SPX} \leq R^k + \Delta)}{\sum_{t=1}^{N} \mathbf{1}(R^k - \Delta < R_t^{SPX} \leq R^k + \Delta)}$$

where $M_{t,t+1}^{daily}(\phi)$ is the time series of the recovered stochastic discount factor evaluated on overlapping one-month intervals sampled at a daily frequency. Considering that risk-free returns are on average small over monthly intervals, it is straighforward to obtain a risk-neutral distribution of returns based on the recovered SDF:

$$f^{MOD}(R;\phi) = f(R)M^{MOD}(R;\phi)$$

### 5.2.2 Empirical Risk-Neutral Density

The risk-neutral density can also be estimated directly from the options data without relying on the recovered SDF. Inspired by the approach of Gatheral and Jacquier (2014), we assume that, on a given date $t$, the implied volatility of the S&P 500 index options with maturity $\tau$ can be represented by the stochastic volatility inspired (SVI) parameterization:

$$w_{t,\tau}(k) = a + b\{\rho(k-m) + \sqrt{(k-m)^2 + \sigma^2}\}$$

where $\sigma^{IV}_{t,\tau}(k) = \sqrt{\frac{w_{t,\tau}(k)}{\tau}}$ is the implied volatility and $k = \log\left(R\frac{S_t}{F_{t,t+\tau}}\right)$ is the normalized strike price written as a function of the return $R$.

We calibrate the five parameters of this surface $(a, b, \rho, m, \sigma)$ for each combination of day and maturity by minimizing the implied volatility mean squared error. The fitted volatilities are subsequently interpolated across maturities to find the one-month values. Given a parametric form $w_t(k)$ for the one-month implied volatility surface, it follows from Breeden and Litzenberger (1978) that the risk-neutral return density is given by

$$f^{SVI}_t(R) = \frac{1}{R\sqrt{w_t(k)}}\left[\left(1 - \frac{kw'_t(k)}{2w_t(k)}\right)^2 - \frac{w'_t(k)^2}{4}\left(\frac{1}{w_t(k)} + \frac{1}{4}\right) + \frac{w''_t(k)}{2}\right]\mathcal{N}\left(\frac{-k}{\sqrt{w_t(k)}} + \frac{\sqrt{w_t(k)}}{2}\right)$$

where $\mathcal{N}$ is the standard normal density. The unconditional distribution $f^{SVI}(R)$ is obtained by simply taking the average of $f^{SVI}_t(R)$ across all dates.

### 5.2.3 Decomposing the Equity Premium

We now compare the data-based characterization of the equity premium with the SDF recovered from the power security returns, which is model-based. Let $f(R)$ and $f^*(R)$ denote the average physical and risk-neutral densities, respectively. Following Beason and Schreindorfer (2022), $EP(x)$ is defined as the relative contribution of returns up to $x$ to the total equity premium, i.e.,

$$EP(x) = \frac{\int_0^x R[f(R) - f^*(R)]dR}{\int_0^\infty R[f(R) - f^*(R)]dR}$$

When inferring the risk-neutral probabilities from the option prices fitted according to the (data-based) SVI parameterization, this expression is given by

$$EP^{DATA}(x) = \frac{\int_0^x R[f(R) - f^{SVI}(R)]dR}{\int_0^\infty R[f(R) - f^{SVI}(R)]dR}$$

Alternatively, if the risk-neutral probabilities are extracted from the (model-based) reciprocal power security return series, we have

$$EP^{MOD}(x;\phi) = \frac{\int_0^x R[f(R) - f^{MOD}(R;\phi)]dR}{\int_0^\infty R[f(R) - f^{MOD}(R;\phi)]dR}$$

### 5.2.4 Empirical Results

Figure 5 presents the empirical results. The black line represents the data. The other lines are based on the SDF recovered from the power security returns according to our model, for different values of $\phi$. The model-based results are not very sensitive to the value of $\phi$. When compared to the data-based $EP$ profile, the model-based results place a bit too much weight on the returns below -20%, and not enough weight on positive returns, suggesting they overestimate the negative skewness in the data and underestimate the kurtosis. However, these deviations are an order of magnitude smaller compared to those for prominent consumption-based models documented in Beason and Schreindorfer (2022).

We conclude that the SDF estimated using the power security returns not only exhibits plausible time-series properties, but also provides a good description of the data when judged by how different parts of the return distribution contribute to the index equity premium.

## 6 Conclusion

We generalize Ross (2015) Recovery Theory by allowing for a persistent non-stationary component in asset prices. This allows for growing economies with nonstationary state variables such as aggregate consumption or the S&P 500. By imposing path-independence, we are able to extract the pricing kernel from long-term power securities. We synthesize these securities using a novel dataset of long-term (up to 10 years) options on the S&P 500 index. The stochastic discount factor is then given by the reciprocal gross return on a long-term security that pays a power of the index.

Our estimate of the stochastic discount factor has plausible time-series properties and is able to explain the equity premium, the long-bond premium, and a negative variance premium. It generates reasonable predictions for return premia on stocks, bonds, and options. The decomposition of the index equity premium across the return state space is mostly consistent with the data for our estimate of the stochastic discount factor, unlike stochastic discount factors in existing equilibrium models (Beason and Schreindorfer, 2020).

# References

Aït-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, *Journal of Finance* 53, 499–547.

Aït-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric Risk Management and Implied Risk Aversion, *Journal of Econometrics* 94, 9–51.

Alvarez, Fernando, and Urban J. Jermann, 2005, Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth, *Econometrica* 73, 1977–2016.

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies* 16, 527–566.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and Derivative-Security Valuation, *Journal of Financial Economics* 55, 205–238.

Barone-Adesi, Giovanni, Nicola Fusari, Antonietta Mira, and Carlo Sala, 2020, Option Market Trading Activity and the Estimation of the Pricing Kernel: A Bayesian Approach, *Journal of Econometrics* 216, 430–449.

Beason, Tyler, and David Schreindorfer, 2020, On Sources of Risk Premia in Representative Agent Models, *Working Paper, Arizona State University.*

Beason, Tyler, and David Schreindorfer, 2022, Dissecting the Equity Premium, *Journal of Political Economy* 130, 2203–2222.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.

Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman, 2016, Misspecified Recovery, *Journal of Finance* 71, 2493–2544.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of State-Contingent Claims Implicit in Option Prices, *Journal of Business* 51, 621–651.

Brennan, Michael J., 1979, The Pricing of Contingent Claims in Discrete Time Models, *Journal of Finance* 34, 53–68.

Chernov, Mikhail, 2003, Empirical Reverse Engineering of the Pricing Kernel, *Journal of Econometrics* 116, 329–364.

Christoffersen, Peter, Steven Heston, and Kris Jacobs, 2013, Capturing Option Anomalies with a Variance-Dependent Pricing Kernel, *Review of Financial Studies* 26, 1963–2006.

Cox, John C., Jonathan C. Ingersoll, and Stephen A. Ross, 1985, A Theory of the Term Structure of Interest Rates, *Econometrica* 53, 385–408.

Cuesdeanu, Horatio, and Jens Carsten Jackwerth, 2018, The Pricing Kernel Puzzle: Survey and Outlook, *Annals of Finance* 14, 289–329.

Epstein, Larry, and Stan Zin, 1989, Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica* 57, 937–969.

Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis, *Journal of Political Economy* 99, 263–286.

Gatheral, Jim, and Antoine Jacquier, 2014, Arbitrage-free SVI volatility surfaces, *Quantitative Finance* 14, 59–71.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2017, What is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models, *Review of Financial Studies* 30, 442–504.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars Peter, and José A. Scheinkman, 2009, Long-Term Risk: An Operator Approach, *Econometrica* 77, 177–234.

Heston, Steven L., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* 6, 327–343.

Heston, Steven L., 2004, Option Valuation with Infinitely Divisible Distributions, *Quantitative Finance* 4, 515–524.

Jackwerth, Jens Carsten, 2000, Recovering Risk Aversion from Option Prices and Realized Returns, *Review of Financial Studies* 13, 433–451.

Jackwerth, Jens Carsten, and Mark Rubinstein, 1996, Recovering Probability Distributions from Option Prices, *Journal of Finance* 51, 1611–1631.

Jensen, Christian Skov, David Lando, and Lasse Heje Pedersen, 2019, Generalized recovery, *Journal of Financial Economics* 133, 154–174.

Kreps, D., and E. Porteus, 1978, Temporal Resolution of Uncertainty and Dynamic Choice Theory, *Econometrica* 46, 185–200.

Linn, Matthew, Sophie Shive, and Tyler Shumway, 2018, Pricing Kernel Monotonicity and Conditional Information, *Review of Financial Studies* 31, 493–531.

Long, John B., 1990, The Numeraire Portfolio, *Journal of Financial Economics* 26, 29–69.

Merton, Robert C., 1973a, An Intertemporal Capital Asset Pricing Model, *Econometrica* 41, 867–887.

Merton, Robert C., 1973b, Theory of Rational Option Pricing, *The Bell Journal of Economics and Management Science* 4, 141–183.

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

Rosenberg, Joshua V., and Robert F. Engle, 2002, Empirical Pricing Kernels, *Journal of Financial Economics* 64, 341–372.

Ross, Steve, 2015, The Recovery Theorem, *Journal of Finance* 70, 615–648.

Rubinstein, Mark, 1976, The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics* 7, 407–425.

Song, Zhaogang, and Dacheng Xiu, 2016, A Tale of Two Option Markets: Pricing Kernels and Volatility Risk, *Journal of Econometrics* 190, 176–196.

Table 1: Monthly Returns. Various Maturities.

| | $m = 6$ | $m = 36$ | $m = 108$ | $m = 120$ |
|---|---|---|---|---|
| **Panel A: Forward** | | | | |
| Mean | 1.00835 | 1.00855 | 1.00937 | 1.00950 |
| SD | 0.04644 | 0.04719 | 0.05086 | 0.05184 |
| Skewness | -0.55305 | -0.50992 | -0.53818 | -0.55828 |
| Kurtosis | 3.82992 | 3.52937 | 3.74062 | 3.85239 |
| | **Correlation Matrix** | | | |
| | 1.00000 | 0.99682 | 0.98997 | 0.98750 |
| | 0.99682 | 1.00000 | 0.99354 | 0.99084 |
| | 0.98997 | 0.99354 | 1.00000 | 0.99960 |
| | 0.98750 | 0.99084 | 0.99960 | 1.00000 |
| **Panel B: Bond** | | | | |
| Mean | 1.00108 | 1.00188 | 1.00359 | 1.00376 |
| SD | 0.00148 | 0.00648 | 0.02231 | 0.02509 |
| Skewness | 1.72295 | 0.30292 | 0.33825 | 0.39732 |
| Kurtosis | 5.77899 | 5.11199 | 4.17194 | 4.41874 |
| | **Correlation Matrix** | | | |
| | 1.00000 | 0.37049 | 0.15267 | 0.13784 |
| | 0.37049 | 1.00000 | 0.78732 | 0.75106 |
| | 0.15267 | 0.78732 | 1.00000 | 0.99752 |
| | 0.13784 | 0.75106 | 0.99752 | 1.00000 |
| **Panel C: Variance Swap** | | | | |
| Mean | -0.00120 | -0.00067 | -0.00121 | -0.00113 |
| SD | 0.01238 | 0.02340 | 0.03374 | 0.03526 |
| Skewness | 5.26050 | 3.66424 | 3.64781 | 3.01674 |
| Kurtosis | 42.77651 | 27.59019 | 30.64465 | 24.23998 |
| | **Correlation Matrix** | | | |
| | 1.00000 | 0.91386 | 0.79671 | 0.74242 |
| | 0.91386 | 1.00000 | 0.91665 | 0.88240 |
| | 0.79671 | 0.91665 | 1.00000 | 0.97593 |
| | 0.74242 | 0.88240 | 0.97593 | 1.00000 |

Notes: Panel A reports the returns on the forward contract, Panel B the returns on the bond, and Panel C the gains on the variance swap. In each Panel, we report descriptive statistics for the time series with maturities of 6 months, 3 years, 9 years, and 10 years, as well as the correlations across maturities. The sample is from January 2007 to December 2023.

Table 2: Descriptive Statistics for Power Security Returns. Various Maturities.

| Maturity | $\phi$ | Mean | SD | Skewness | Kurtosis |
|---|---|---|---|---|---|
| | 0.00 | 1.00108 | 0.00148 | 1.72295 | 5.77899 |
| | 0.25 | 1.00311 | 0.01210 | -0.64650 | 4.28295 |
| | 0.50 | 1.00500 | 0.02388 | -0.62379 | 4.12979 |
| | 0.75 | 1.00674 | 0.03532 | -0.58820 | 3.96993 |
| 6 | 1.00 | 1.00835 | 0.04644 | -0.55305 | 3.82992 |
| | 0.00 | 1.00188 | 0.00648 | 0.30292 | 5.11199 |
| | 0.25 | 1.00371 | 0.01316 | -0.55790 | 3.96826 |
| | 0.50 | 1.00545 | 0.02492 | -0.58519 | 3.87555 |
| | 0.75 | 1.00706 | 0.03635 | -0.55514 | 3.70246 |
| 36 | 1.00 | 1.00855 | 0.04719 | -0.50992 | 3.52937 |
| | 0.00 | 1.00359 | 0.02231 | 0.33825 | 4.17194 |
| | 0.25 | 1.00519 | 0.02127 | -0.70488 | 5.74571 |
| | 0.50 | 1.00671 | 0.02907 | -0.84842 | 5.91214 |
| | 0.75 | 1.00810 | 0.03977 | -0.67331 | 4.53963 |
| 108 | 1.00 | 1.00937 | 0.05086 | -0.53818 | 3.74062 |
| | 0.00 | 1.00376 | 0.02509 | 0.39732 | 4.41874 |
| | 0.25 | 1.00533 | 0.02316 | -0.68444 | 5.94570 |
| | 0.50 | 1.00683 | 0.03020 | -0.88737 | 6.26065 |
| | 0.75 | 1.00822 | 0.04070 | -0.70253 | 4.73689 |
| 120 | 1.00 | 1.00950 | 0.05184 | -0.55828 | 3.85239 |

Notes: We report descriptive statistics for monthly power security returns that pay out $(S_T)^{\phi}$, computed using different values of $\phi$ and different maturities. The sample is from January 2007 to December 2023.

Table 3: Return Correlations

Panel A: Correlations Between Power Security Returns for Different Maturities

| | $m = 6$ | $m = 36$ | $m = 108$ | $m = 120$ |
|---|---|---|---|---|
| $m = 6$ | 1.00000 | 0.98961 | 0.91785 | 0.90210 |
| $m = 36$ | 0.98961 | 1.00000 | 0.94494 | 0.92973 |
| $m = 108$ | 0.91785 | 0.94494 | 1.00000 | 0.99819 |
| $m = 120$ | 0.90210 | 0.92973 | 0.99819 | 1.00000 |

Panel B: Correlations Between Security Returns

| | | | | Power Security | | |
|---|---|---|---|---|---|---|
| | Forward | Bond | Variance | $\phi = 0.25$ | $\phi = 0.5$ | $\phi = 0.75$ |
| Forward | 1.00000 | -0.06678 | -0.55292 | 0.57988 | 0.90810 | 0.98777 |
| Bond | -0.06678 | 1.00000 | 0.01472 | 0.76581 | 0.34059 | 0.07666 |
| Variance | -0.55292 | 0.01472 | 1.00000 | -0.43308 | -0.60520 | -0.59904 |
| $\phi = 0.25$ | 0.57988 | 0.76581 | -0.43308 | 1.00000 | 0.86443 | 0.69526 |
| $\phi = 0.5$ | 0.90810 | 0.34059 | -0.60520 | 0.86443 | 1.00000 | 0.96174 |
| $\phi = 0.75$ | 0.98777 | 0.07666 | -0.59904 | 0.69526 | 0.96174 | 1.00000 |

Notes: Panel A reports correlations between six-month, 3-year, 9-year, and 10-year power security returns with $\phi = 0.5$. Panel B reports correlations between the power security returns with multiple $\phi$ values, the forward return, the bond return, and the variance swap gain, all with 10-year maturity. The sample is from January 2007 to December 2023.

Table 4: Regressing the Power Security on the Forward, Bond, and Variance Swap

| | $\phi = 0.0$ | $\phi = 0.25$ | $\phi = 0.50$ | $\phi = 0.75$ | $\phi = 1.0$ |
|---|---|---|---|---|---|
| $\log R^{FWD}$ | 0.000 | 0.249*** | 0.499*** | 0.751*** | 1.000*** |
| | (0.000) | (0.001) | (0.001) | (0.001) | (0.000) |
| $\log R^{B}$ | 1.000*** | 0.747*** | 0.487*** | 0.233*** | −0.000 |
| | (0.000) | (0.002) | (0.003) | (0.002) | (0.000) |
| $R^{VAR}$ | −0.000 | −0.087*** | −0.114*** | −0.082*** | −0.000 |
| | (0.000) | (0.004) | (0.006) | (0.004) | (0.000) |
| Constant | 0.000 | 0.0004*** | 0.0005*** | 0.0004*** | 0.000** |
| | (0.000) | (0.0001) | (0.0001) | (0.0001) | (0.000) |
| Observations | 203 | 203 | 203 | 203 | 203 |
| Adjusted $R^2$ | 1.000 | 0.999 | 0.999 | 1.000 | 1.000 |
| *Note:* | | | | | *p<0.1; **p<0.05; ***p<0.01 |

Notes: We regress the power security log returns computed using different values of $\phi$ on the forward log return, bond log return, and variance swap gain, all with 10-year maturity. The sample is from January 2007 to December 2023.

Table 5: Stochastic Discount Factors: Descriptive Statistics

| $\phi$ | N | Mean | SD | Skewness | Kurtosis | $\rho(1)$ |
|---|---|---|---|---|---|---|
| 0.00 | 203 | 0.9969 | 0.025 | -0.17 | 4.00 | 0.060 |
| 0.25 | 203 | 0.9952 | 0.023 | 1.01 | 6.97 | -0.052 |
| 0.50 | 203 | 0.9941 | 0.031 | 1.34 | 8.36 | -0.036 |
| 0.75 | 203 | 0.9935 | 0.042 | 1.15 | 6.54 | 0.000 |
| 1.00 | 203 | 0.9933 | 0.053 | 1.00 | 5.15 | 0.020 |

Notes: We report the mean, standard deviation (SD), skewness, kurtosis, and first-order autocorrelation ($\rho(1)$) for the estimated time series of the pricing kernel with different values of $\phi$. The sample is from January 2007 to December 2023.

Figure 1: S&P 500 Return, Realized Variance, and Cumulative Return

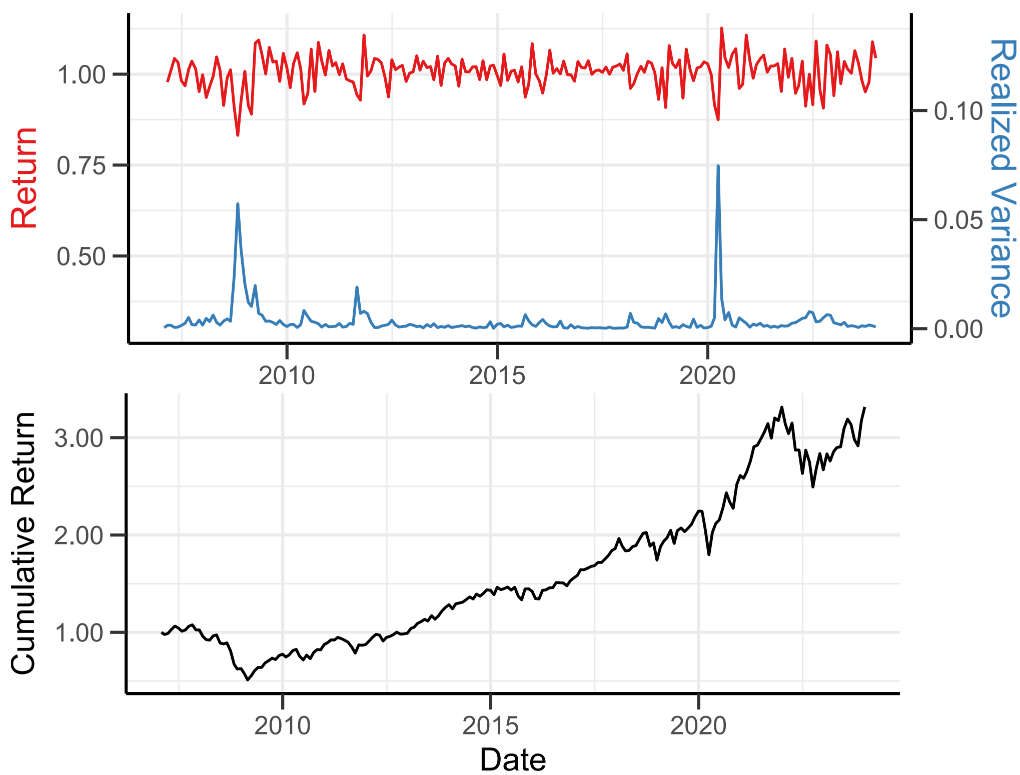


Notes: The top panel plots the S&P 500 return (left axis) and realized variance (right axis) for non-overlapping monthly intervals. The bottom panel plots the cumulative return on the S&P 500. The sample is from January 2007 to December 2023.

Figure 2: Bond, Forward, Variance Swap, and Power Security

Panel A: Bond Return, Forward Return, and Variance Swap Gain.

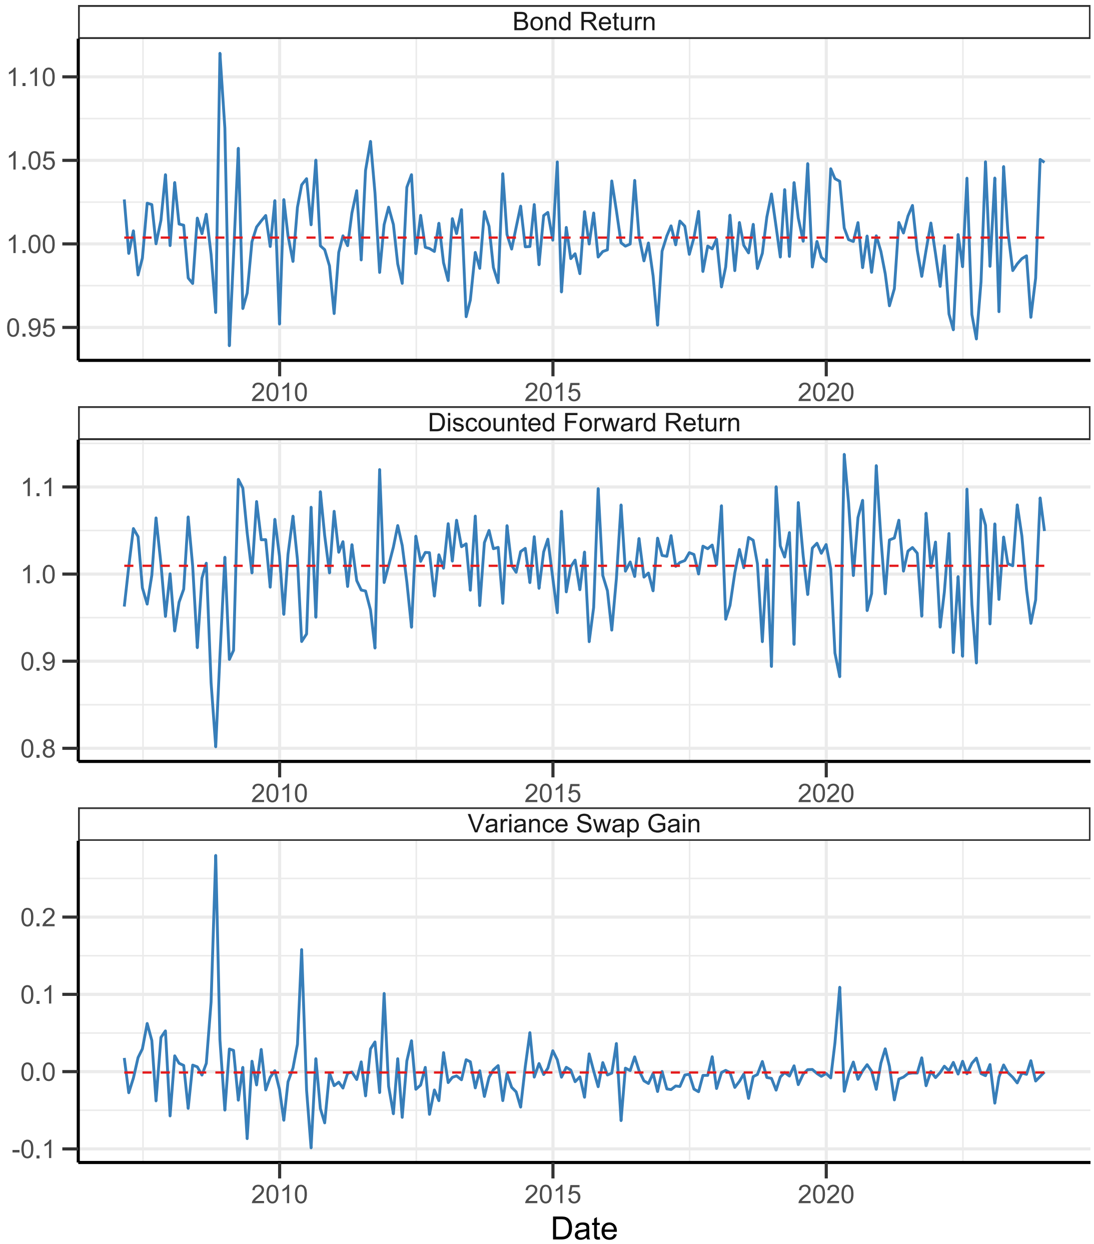

Panel B: Power Security Returns

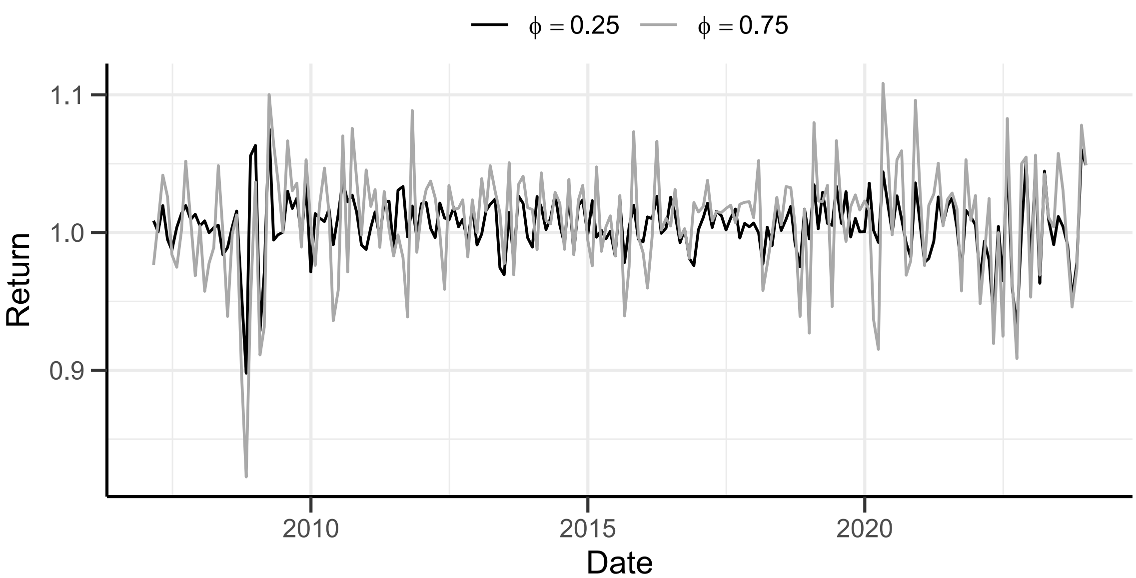


Notes: Panel A plots the forward return, the bond return, and the variance swap gain, all with the 10-year maturity. Panel B plots the 10-year power security returns computed using different values of $\phi$. The sample is from January 2007 to December 2023.

Figure 3: Interpreting the Power Security Return

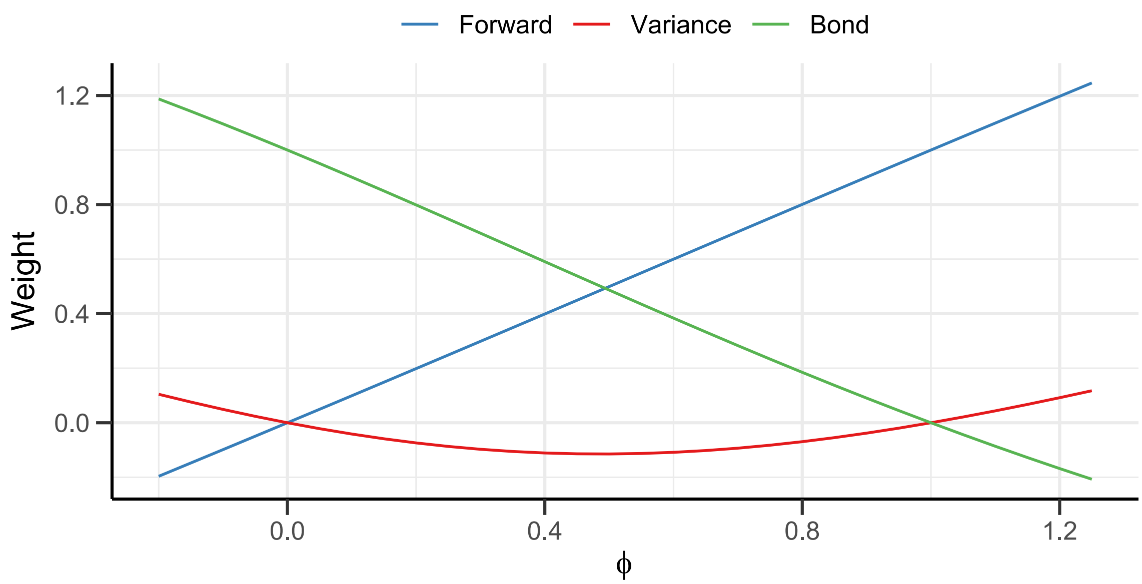


Notes: Following Table 4, we interpret the 10-year power security as a portfolio consisting of a forward, a bond, and a variance swap with the same maturity and show how the portfolio weights vary across a fine grid of $\phi$ values. The sample is from January 2007 to December 2023.

Figure 4: The Time Series of the Stochastic Discount Factor

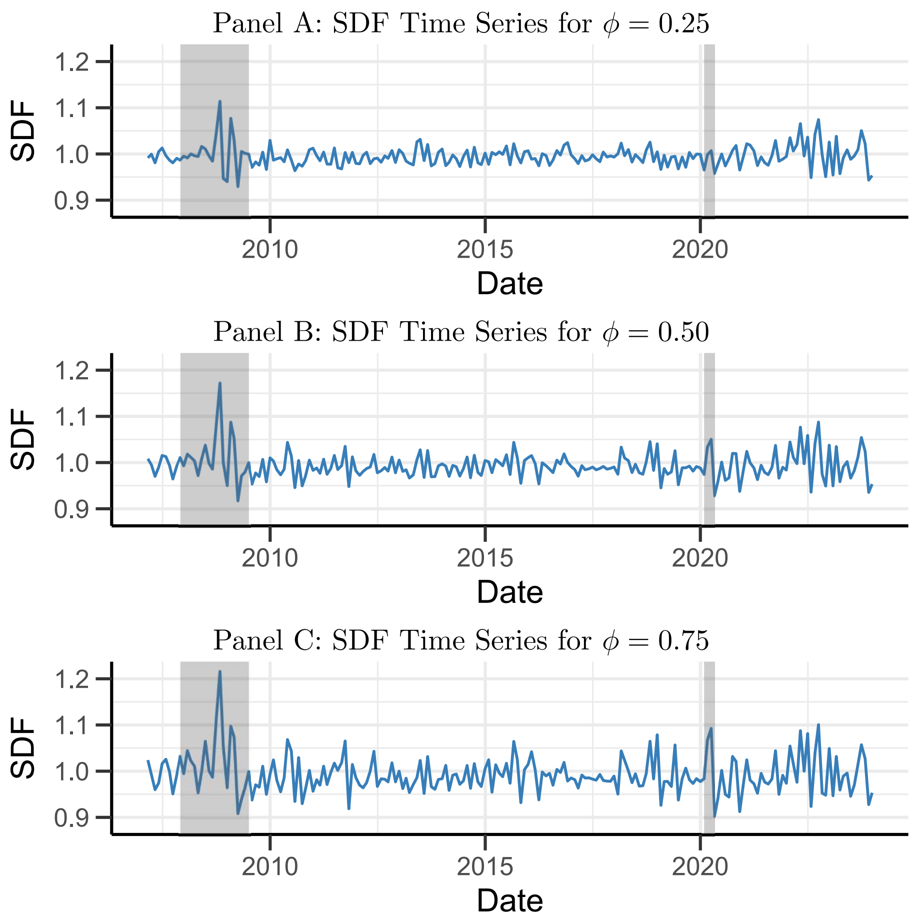


Notes: We plots the time series of the estimated stochastic discount factor for different values of $\phi$. The sample is from January 2007 to December 2023.

Figure 5: Decomposing the Equity Premium

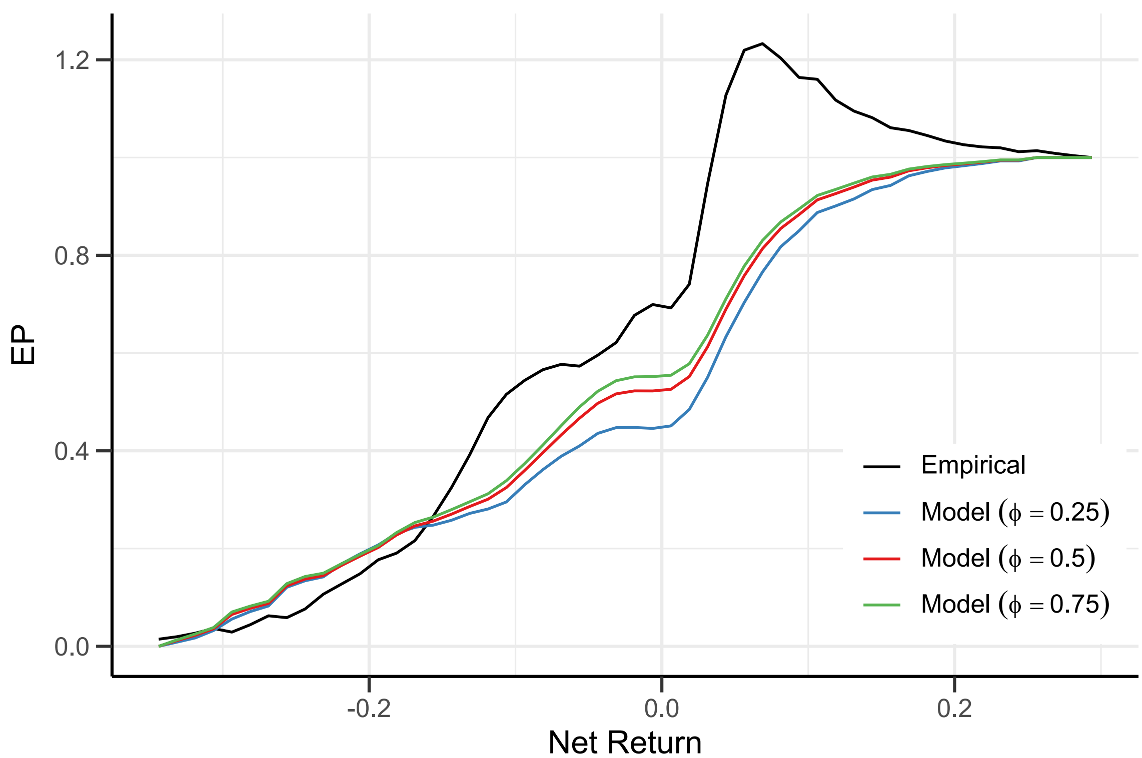


Notes: We decompose the equity premium following Beason and Schreindorfer (2022). On the horizontal axis are net returns. The figure shows how much of the overall equity premium is due to different return levels on average over the sample. The black line shows the result when the risk-neutral density is extracted directly from option prices, while the other lines are based on the SDF recovered from the power security returns according to our model with different $\phi$ values. The sample is from January 2007 to December 2023.