Arbitrage-Based Recovery

Ferenc Horvath*

City University of Hong Kong

30 June, 2021

Abstract

We develop a novel recovery theorem based on no-arbitrage principles. Our Arbitrage-Based Recovery Theorem does not require assuming time homogeneity of either the physical probabilities, the Arrow-Debreu prices, or the stochastic discount factor; and it requires the observation of Arrow-Debreu prices only for one single maturity. We perform several different density tests and mean prediction tests using 25 years of S&P 500 options data, and we find evidence that our method can correctly recover the probability distribution of the S&P 500 index level on a monthly horizon.

JEL classification: G1, G12, G13.
Keywords: recovery theorem, physical probabilities, transition state prices, no-arbitrage.

*City University of Hong Kong, Department of Economics and Finance, 83 Tat Chee Avenue, Kowloon Tong, Kowloon, Hong Kong. Tel.: +852 3442 7630. E-mail address: f.horvath@cityu.edu.hk

1 Introduction

Extracting information about the physical probabilities of future events from market data has been the focus of finance research for several decades. A huge variety of applications in financial economics relies on knowledge of physical probabilities, such as portfolio choice, risk management, and asset pricing, just to name a few. Market practice usually hinges upon estimating physical probabilities from historical data. Nevertheless, estimates are only as reliable as good of a representative the historical data are. Therefore, extracting physical probabilities from real-time market data in a forward-looking manner (as opposed to estimating them from backward-looking historical data) would be of paramount importance.

However, since the seminal works of Black and Scholes (1973) and Merton (1973) (and even earlier, of Bachelier (1900)), we know that asset prices, per se, depend only on risk-neutral probabilities, but not on physical probabilities. Therefore, extracting information about physical probabilities from asset pricing data is not possible without imposing further assumptions. More than four decades had passed until the first such set of assumptions was stated and the corresponding physical probability recovery theory developed by Ross (2015). In a one-period finite-state framework Ross demonstrates that if we are able to observe the entire Arrow-Debreu price matrix and the stochastic discount factor (SDF) is the product of a constant and a fraction, the numerator of which depends only on the arrival state and the denominator only on the initial state, then the physical probability measure can be recovered. Since in reality we can observe only one row of the Arrow-Debreu price matrix (and not the entire matrix), Ross extends his recovery theorem into a multiperiod Markovian framework, where he assumes that Arrow-Debreu prices are time-homogeneous, and the one-period SDF takes the same functional form as in his one-period model. Under these assumptions, if we can observe the transition state prices for as many maturities as the number of possible states, the entire Arrow-Debreu matrix can be recovered[1] and, as a consequence, one-period physical transition probabilities can also be recovered.

Ross's recovery theorem has initiated a lively academic debate both in the theoretical and the empirical finance literature. While Hansen and Scheinkman (2009) show that the stochastic discount factor process can be expressed as the product of Ross's SDF and a martingale, Borovička, Hansen, and Scheinkman (2016) prove that Ross's approach recovers the true physical probabilities if and only if the martingale

[1] As Jensen, Lando, and Pedersen (2019) point out, the recovered Arrow-Debreu matrix is almost surely unique. I.e., the set of parameters for which there exists a continuum of recovered Arrow-Debreu matrices (instead of a unique recovered Arrow-Debreu matrix) has a measure of zero.

component of the SDF process is constant. They also provide several illustrative examples where the martingale component of the SDF is not constant, including the long-run risk model of Bansal and Yaron (2004). Using options on the 30-year Treasury bond futures, Bakshi, Chabi-Yo, and Gao (2018) reject the implication that the martingale component of the SDF is constant.

Besides assuming a constant SDF martingale component, another strong assumption of Ross's recovery theorem is that transition state prices are time-homogeneous. Jensen, Lando, and Pedersen (2019) develop the *Generalized Recovery Theorem*, where they relax this assumption. They show that as long as we are able to observe the transition state prices for at least as many maturities as the number of possible states, physical probabilities can almost surely be recovered, without assuming time-homogeneity.

Jensen, Lando, and Pedersen (2019) also demonstrate when (and how) *Generalized Recovery* can be implemented even if the number of possible states grows over time. Standard examples for such a case include the models of Mehra and Prescott (1985) and of Cox, Ross, and Rubinstein (1979). Basically, one needs to make use of the fact that the SDF can be expressed as a function (with a *fixed* number of parameters) of the state variable. Recovery then means recovering the parameter values. Then, after calculating the SDF values, recovering the physical probabilities is straightforward. In this structured version of the *Generalized Recovery Theorem*, one needs to observe the transition state prices for at least as many maturities as the number of unknown parameters.

An unattractive feature of the currently available recovery theorems is that they require observing the transition state prices for several different maturities. Even in the simplest case when the only state variable is the asset price, one needs to observe the transition state prices for at least as many different maturities as many possible values the asset price can take. For example, if one is interested in the physical probabilities over a one-month horizon and one can observe the monthly transition state prices for up to one year of maturity, then (assuming non-overlapping observation periods) this allows for a model where the asset price can take only twelve different values. As Ross (2015) notes, such a coarse grid leads to poorly discretized transition state prices and eventually poorly discretized recovered probabilities. Audrino, Huitema, and Ludwig (2019) use overlapping observation periods and transition state prices of 1/10 of a month, but without additional restrictions, the recovered one-month transition state prices still exhibit counterintuitive features, as it is demonstrated by Jackwerth and Menner (2020). An adapted version of the Generalized Recovery approach of Jensen, Lando, and Pedersen (2019) suggests a way out: if the SDF can be written as a known function (with a small number of unknown parameters) of the state variables, then it is enough to observe the transition state prices for only as many different maturities as the number of parameters (which is, ideally, much lower than the number of possible values of the state variable). There is, however, a trade-off: the modeler has to exogenously provide the functional form of the SDF. The validity of the recovered physical probabilities will then also depend on the validity of the assumed SDF functional form. Furthermore, one still needs to observe the transition state prices for several different maturities, if the SDF has more than one parameter.

Jackwerth and Menner (2020) perform a thorough empirical analysis to assess whether the recovered physical probabilities are indeed equal to the true physical probabilities. Using more than 30 years of S&P 500 European options monthly observations, they recover the 1-month physical probability distributions of the S&P 500 index level via different versions of the recovery theorem. Then, based on several different density tests and mean- and variance prediction tests, Jackwerth and Menner reject the hypothesis that realized S&P 500 index values are drawn from the recovered physical probability distributions.

The empirical analysis of Jackwerth and Menner (2020) also points out that applying Ross's recovery theorem to reconstruct the entire Arrow-Debreu matrix can lead to a recovered transition state price matrix with highly counterintuitive features. E.g., even though one would expect high recovered state prices on the main diagonal of the Arrow-Debreu price matrix, the recovered Arrow-Debreu matrix of Jackwerth and Menner tends to have high state prices in states which are far off-diagonal. Furthermore, the model-implied risk-free rates are often negative or their magnitude is as high as several hundred percent.

An aspect of recovery theorems, which has so far been unexplored in the literature, is whether they allow for arbitrate opportunities or not. Seemingly, it is enough to assume non-negative Arrow-Debreu prices to exclude arbitrage opportunities. However, as we demonstrate in Section 3, this is only a necessary but not a sufficient condition in the realm of recovery theorems. Actually, implementing either Ross's recovery or the generalized recovery does not necessitate thinking about arbitrage opportunities. However, we show that imposing the assumption of no arbitrage, several restrictive assumptions of the other recovery theorems (e.g., time homogeneity of transition state prices) can be relaxed. Furthermore, empirical implementation becomes much easier since we do not need to observe the Arrow-Debreu prices for several different maturities.

From a theoretical perspective, our Arbitrage-Based Recovery Theorem thus provides a solution to the above-mentioned reservations: it does not assume time homogeneity, it requires observing the transition state prices only for one single maturity, and it explicitly excludes arbitrage opportunities. Since the defense of a theory resides in its empirical validity, we use 25 years of S&P 500 options data to recover the one-month physical probabilities of the S&P 500 index level changes implied by our recovery theorem. Then, we perform several density tests (Berkowitz test, two versions of the Knüppel test, Kolmogorov-Smirnov test) and mean prediction tests to assess whether the S&P 500 movements are indeed drawn from the physical probability distribution recovered by our theorem. The answer is affirmative: while Jackwerth and Menner (2020) find (using almost identical data and tests) that empirical data do not back up the recovered probabilities of other recovery theorems, the hypothesis that our theorem recovers the true physical probabilities cannot be rejected by the majority of the same tests. Furthermore, we find that our recovery theorem performs especially well if the market of options written on the S&P 500 index on the observation date is characterized by a "volatility smile" and not by a "volatility skew".

This paper is organized as follows. Section 2 describes the recovery theorems currently available in the literature: Ross's recovery in its single-period and multiperiod forms, and the standard and large-state-space versions of the generalized recovery of Jensen, Lando, and Pedersen (2019). In Section 3 we develop our Arbitrage-Based Recovery Theorem. Section 4 contains our empirical tests, while Section 5 concludes.

2 Evolution of recovery theorems

In this section, we demonstrate the evolution of recovery theorems in the literature from Ross's Recovery Theorem through the Generalized Recovery of Jensen, Lando, and Pedersen (2019) to the Arbitrage-Based Recovery of this paper. We provide a brief summary of the recovery theorems in Table 1.

2.1 Ross's recovery (single period)

The very first recovery theorem was developed by Ross (2015). The framework considers a single-period complete financial market, in which the state variable can take n different values. The $n \times n$ matrix of (one-period) Arrow-Debreu prices is $\boldsymbol{A}(1)$, the elements of which are strictly positive. Ross (2015) then proves that if the (one-period) SDF has the functional form

$$m(1)_{ij} = \delta \frac{h(\theta_j)}{h(\theta_i)}, \tag{1}$$

Table 1: Comparison of recovery theorems

	Ross (2015) (single period)	Ross (2015) (multi-period)	Audrino, Huitema, and Ludwig (2019)	Jackwerth and Menner (2020) (Ross Stable)	Jensen, Lando, and Pedersen (2019)	This paper
Requires observing only one maturity?	✓	✗	✗	✗	✗	✓
Requires observing only the current state (i.e., not "parallel universes")?	✗	✓	✓	✓	✓	✓
Valid without assuming time homogeneity?	✓	✗	✗	✗	✓	✓
Valid without assuming constant Perron-Frobenius eigenvectors and a geometric sequence of Perron-Frobenius eigenvalues?	✓	✓	✓	✓	✗	✓
Does it preclude arbitrage opportunities?	✗	✗	✗	✗	✗	✓

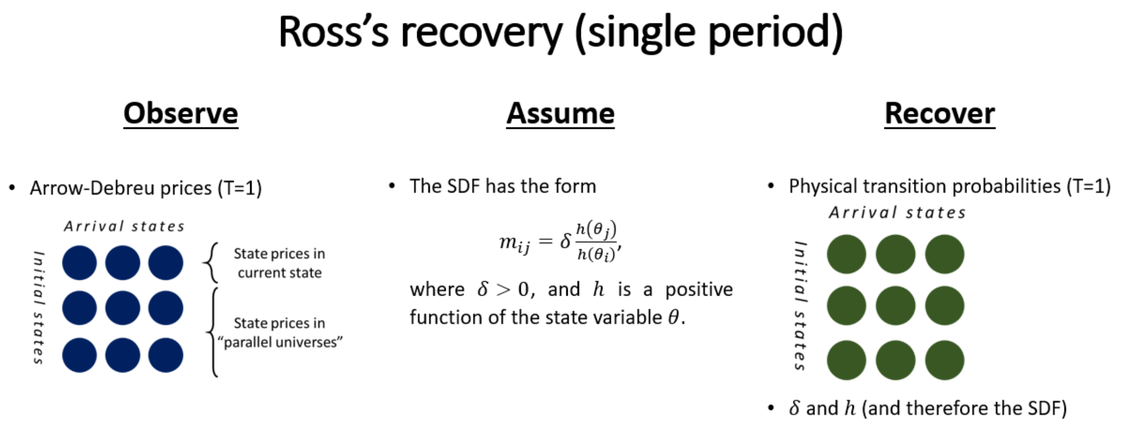

Figure 1: Ross's recovery (single period). One must observe the whole Arrow-Debreu price matrix. Then, assuming that the SDF (going from state i to state j) takes the form $\delta \times h(\theta_j)/h(\theta_i)$, where $\delta > 0$ and h is a positive function of the state variable θ, Ross proves that the physical transition probabilities and the SDF can be recovered.

where δ is a positive constant, θ_i is the state variable in state i, and h is a positive function of the states, then the Arrow-Debreu prices $A(1)_{ij}$ can be decomposed uniquely into the product of the SDF $m(1)_{ij}$ and the (one-period) physical transition probabilities $\pi(1)_{ij}$. We summarize this recovery theorem in Figure 1. The downside of this approach is that it requires observing the entire Arrow-Debreu price matrix. I.e., we need to know the state prices not only when the initial state is the current state, but also when the initial state is different from the current state. Since the state prices of such "parallel universes" cannot be observed, this recovery theorem cannot be implemented in practice.

2.2 Ross's recovery (multiperiod, non-overlapping periods)

To implement his recovery theorem in practice, Ross (2015) extends his single-period model into a multiperiod setup, where the state variable follows a Markov chain on $\{1, \ldots, n\}$. Furthermore, he imposes the additional assumption that the transition state prices are time homogeneous. I.e.,

$$\boldsymbol{A}(T) = \boldsymbol{A}(1)^T, \tag{2}$$

where $\boldsymbol{A}(T)$ is the Arrow-Debreu price matrix corresponding to a maturity of T periods. Then, observing the current transition state prices (i.e., only one row of the Arrow-Debreu price matrix) for n different maturities, one can write down n^2 equations with n^2 unknown variables, and in principle the entire one-period Arrow-Debreu matrix $\boldsymbol{A}(1)$ can be recovered. Afterwards, the single-period recovery theorem can

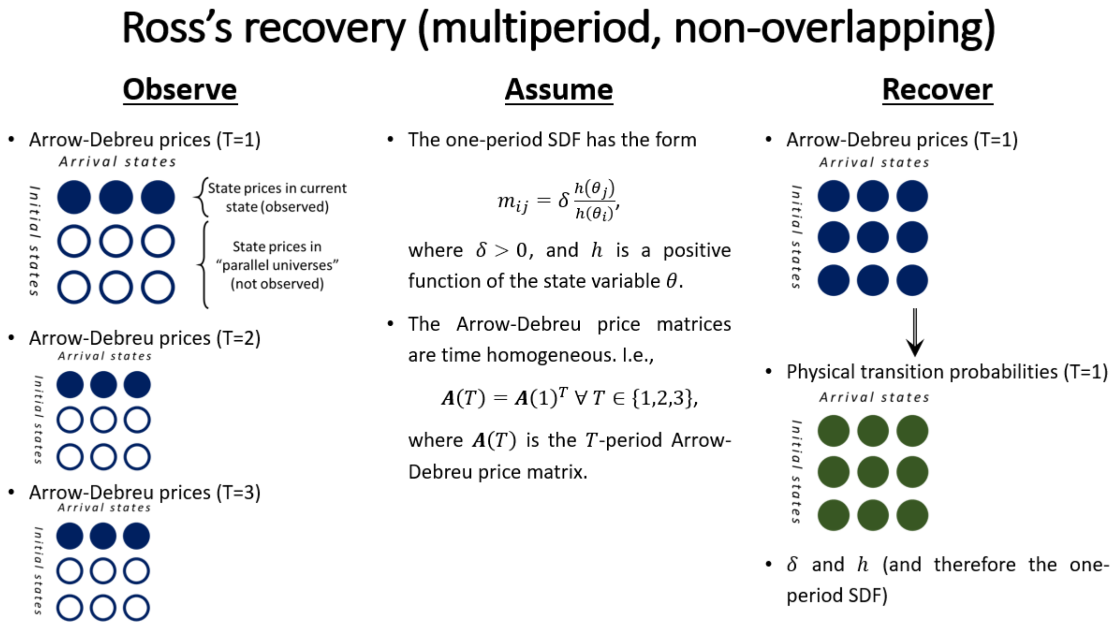

Figure 2: Ross's multiperiod recovery with non-overlapping periods. One must observe the Arrow-Debreu prices (where the current state is the initial state) for as many maturities as many different values the state variable can take. If the Arrow-Debreu price matrices are time homogeneous, then the whole one-period Arrow-Debreu price matrix can be recovered. Then, if the one-period SDF takes the same functional form as assumed in Ross's single-period recovery theory, the (one-period) physical probabilities and the (one-period) SDF itself can also be recovered.

be readily applied, and the one-period Arrow-Debreu prices can be uniquely decomposed into the product of the one-period SDF and the one-period physical transition probabilities. We summarize this version of Ross's recovery theorem in Figure 2. The unappealing feature of this approach is that the assumption of time-homogeneous transition state prices (and consequently time-homogeneous risk-neutral probabilities) is rather restrictive and unlikely to be supported empirically. Moreover, the state variable can take only as many different values as the number of observations, which can lead to poorly discretized state prices and transition probabilities.

2.3 Ross's recovery (multiperiod, overlapping periods)

The problem of poorly discretized transition probabilities due to observing the current state prices for too few maturities can be mitigated by using overlapping observation periods, as pointed out by Audrino, Huitema, and Ludwig (2019). We demonstrate this approach in Figure 3. Since every new observation provides n additional equations, increasing the number of observed maturities enables building a finer grid of possible state prices. Imposing the assumption of time-homogeneous state price matrices, the

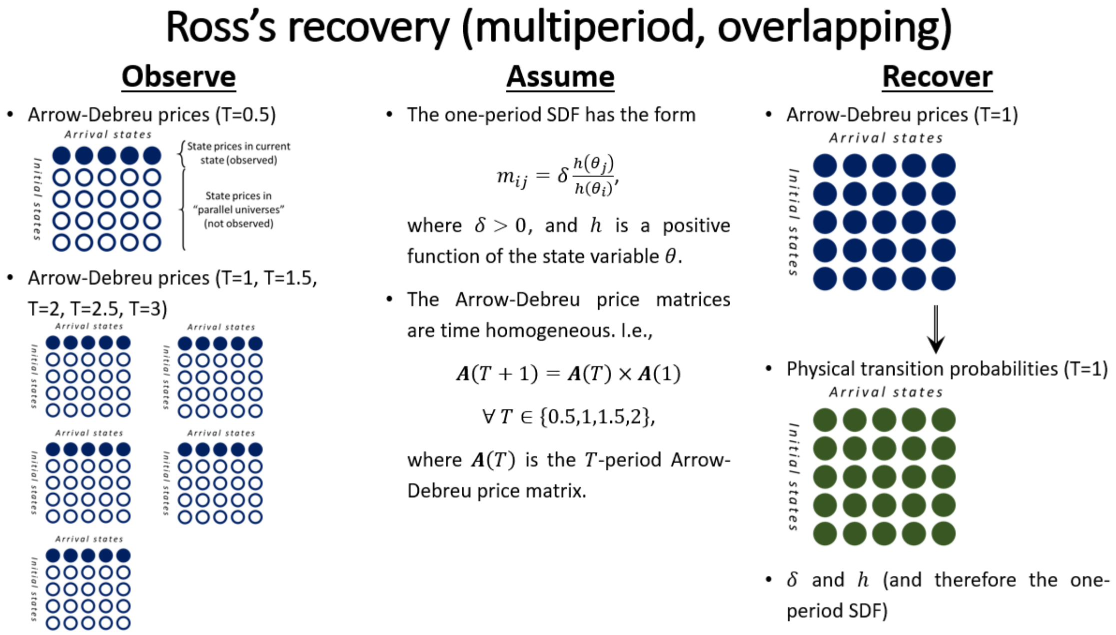

Figure 3: Ross's multiperiod recovery with overlapping periods. One must observe the Arrow-Debreu prices (where the current state is the initial state) for as many maturities as many different values the state variable can take. If the Arrow-Debreu price matrices are time homogeneous, then the whole one-period Arrow-Debreu price matrix can be recovered. Then, if the one-period SDF takes the same functional form as assumed in Ross's single-period recovery theory, the (one-period) physical probabilities and the (one-period) SDF itself can also be recovered. Note that using overlapping periods, one can observe more maturities than with non-overlapping periods and therefore a finer grid can be built.

entire one-period Arrow-Debreu price matrix can be reconstructed. Then, recovering the one-period physical transition probabilities on the finer grid is a straightforward application of Ross's single-period recovery theorem.

Although using overlapping observation periods allows for a finely discretized state price grid, this approach still retains the rather strong assumption of time homogeneity. Furthermore, as Jackwerth and Menner (2020) show using options on the S&P 500 index, the obtained Arrow-Debreu price matrix exhibits several counterintuitive features. These counterintuitive features can be mitigated only by imposing additional constraints while solving for the Arrow-Debreu prices, such as bounding the model-implied one-period risk-free rate from below and above, or imposing unimodality on the rows of the Arrow-Debreu price matrix.

2.4 Generalized Recovery (standard version)

To overcome the limitations imposed by time homogeneity, Jensen, Lando, and Pedersen (2019) develop an alternative approach to recover the physical transition probabilities, which they call *Generalized Recovery* (Figure 4). Similarly to the multiperiod version of Ross's recovery theorem (with non-overlapping observation periods), this approach requires observing the current transition state prices for as many maturities as the number of possible state variable values. The appeal of the Generalized Recovery approach is, however, that it does not assume time homogeneity of either the Arrow-Debreu price matrices or the transition probabilities. Instead, it assumes that the Perron-Frobenius eigenvectors of the Arrow-Debreu price matrices are constant, and their Perron-Frobenius eigenvalues form a geometric sequence. From an economic perspective, this assumption is equivalent to assuming that the SDF with maturity T takes the form

$$m(T)_{ij} = \delta^T \frac{h(\theta_j)}{h(\theta_i)}, \tag{3}$$

where $\delta > 0$ and h is a positive function of the state variable θ.

Under these assumptions, the SDF and the physical transition probabilities (where the current state is the initial state) can be recovered for each observed maturity. Importantly, one cannot recover the whole transition probability matrices, only their row which corresponds to the current state as the initial state. But this is hardly a limitation, since only those actual probabilities are of interest (as opposed to the probabilities of "parallel universes").

An alternative derivation of the Generalized Recovery method can be found in Jackwerth and Menner (2020) under the name "Ross Stable". Although the physical probabilities recovered by the Generalized Recovery method and the Ross Stable approach are identical, the two theories differ in their fundamental assumption. Jackwerth and Menner (2020) retain the assumption of time-homogeneous Arrow-Debreu prices, and they do not make any explicit assumption about the relationship between the SDFs of different maturities. In contrast to this, Jensen, Lando, and Pedersen (2019) explicitly abandon the assumption of time homogeneity, but they assume constant Perron-Frobenius eigenvectors and proportional Perron-Frobenius eigenvalues. I.e., they assume that the T-period SDF takes the form (3).[2]

[2] As Jensen, Lando, and Pedersen (2019) note, in a Ross economy (where Arrow-Debreu prices are time homogeneous) their recovered probabilities are identical to the probabilities recovered by Ross. This observation is in strong relationship with the seeming equivalence of the Ross Stable method and the Generalized Recovery approach. Concretely, Jackwerth and Menner (2020) show that if the Arrow-Debreu price matrices are time homogeneous, then the Perron-Frobenius eigenvectors are constant and the Perron-Frobenius eigenvalues are proportional to each other.

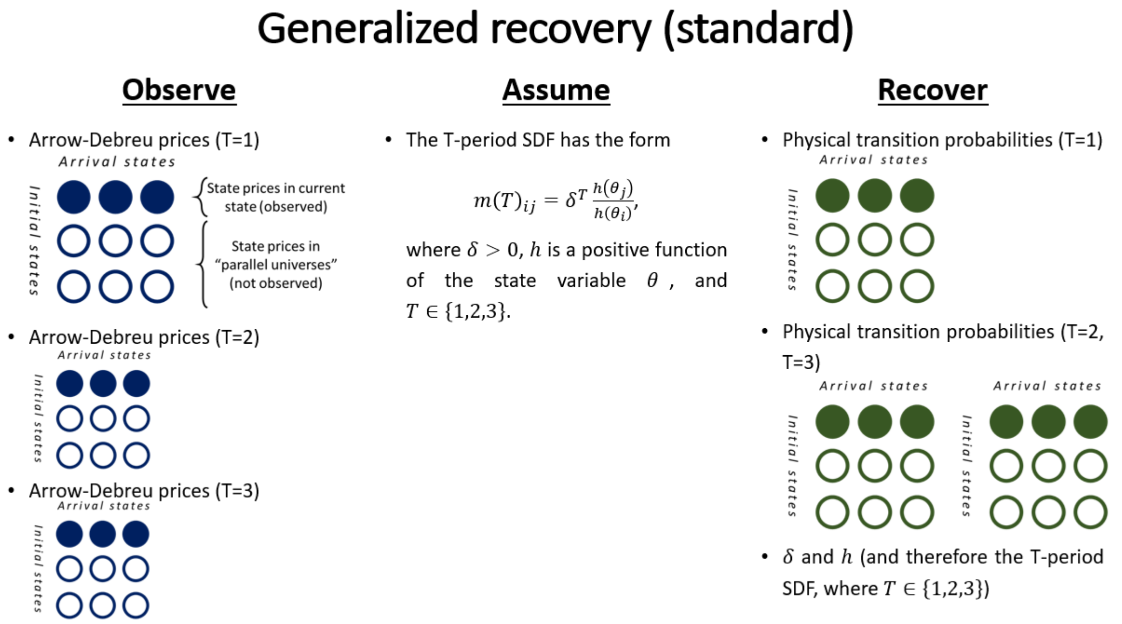

Figure 4: Generalized recovery (standard version). One must observe the Arrow-Debreu prices (where the current state is the initial state) for as many maturities as many different values the state variable can take. If the T-period SDF takes the form $m(T)_{ij} = \delta^T \times h(\theta_j)/h(\theta_i)$, where $\delta > 0$ and h is a positive function of the state variable θ, then the physical transition probabilities (where the current state is the initial state) and the SDF can be recovered for each observed maturity.

2.5 Generalized recovery (large state space version)

When the state variable can take a large number of different values, observing current state prices for such a large number of different maturities might be unfeasible. An alternative version of the generalized recovery of Jensen, Lando, and Pedersen (2019) suggests a solution (Figure 5). Concretely, they show that if the T-period SDF takes the functional form (3) and h is a known positive function of the state variable with a small number of unknown parameters, then by observing the current Arrow-Debreu prices (where the current state is the initial state) for at least by one more maturity than the number of unknown parameters of h, the SDFs can be recovered. As a consequence, the current physical transition probabilities (where the current state is the initial state) can also be recovered for each observed maturity.

This version of the generalized recovery theory allows for a large state space and a finely discretized grid. Furthermore, although one still needs to observe the current state prices for several different maturities, the number of these maturities needs to be only as large as the number of unknown parameters in the SDFs. The drawback of this approach is that one needs to exogenously impose a particular functional form on the SDF (ideally with a small number of parameters), and the choice of this functional

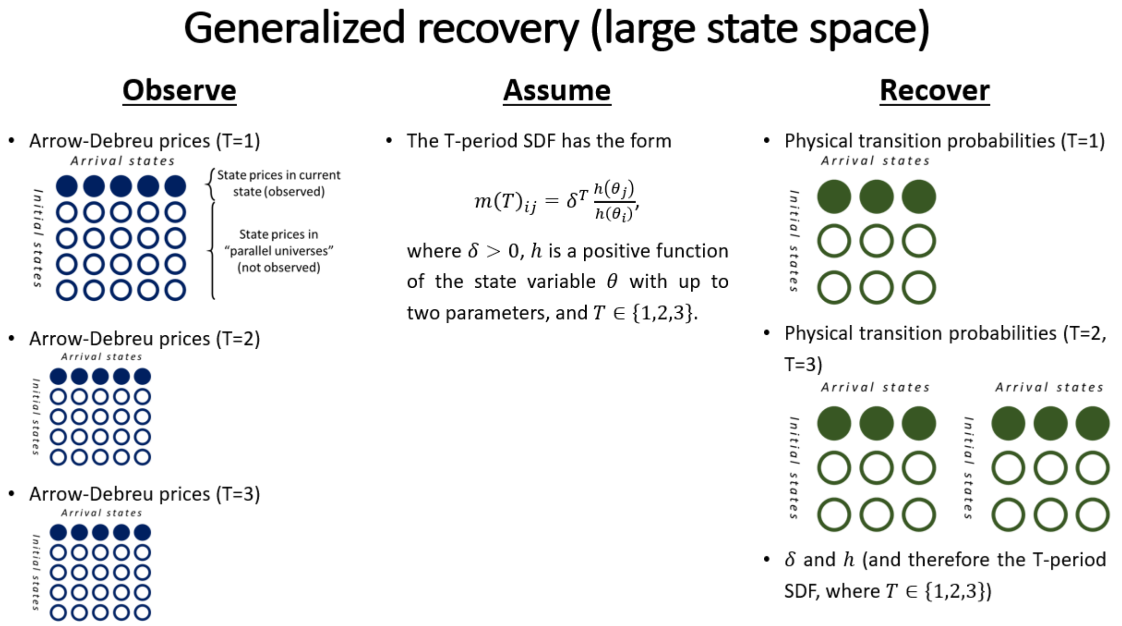

Figure 5: Generalized recovery (large state space version). The T-period SDF is assumed to take the form $m(T)_{ij} = \delta^T \times h(\theta_j)/h(\theta_i)$, where $\delta > 0$ and h is a known positive function of the state variable θ with a small number of unknown parameters. One must observe the Arrow-Debreu prices (where the current state is the initial state) for at least by one more maturity as the number of the unknown parameters of h. Then, the physical transition probabilities (where the current state is the initial state) and the SDF can be recovered for each observed maturity.

form eventually also affects the recovered physical probabilities.

3 Arbitrage-based recovery

In this paper, we develop a novel recovery theorem which requires neither the assumption of time homogeneity nor the observation of Arrow-Debreu prices for several different maturities. Applying our version of the recovery theorem, we need to observe the Arrow-Debreu prices only for the current state as initial state and only for one single maturity, and from those observations we can recover the physical transition probabilities corresponding to that maturity. Our model can be applied to recover the physical probabilities of the price movements of any assets, as long as the state space is spanned by the price of that asset.[3] A schematic demonstration of our recovery theory can be found in Figure 6.

Consider a single-period complete financial market. For the sake of concreteness, we consider here the market for a stock, but our approach can readily be applied for

[3]If there are other state variables besides the asset price, then our approach can be used to recover the marginal probability distribution followed by the asset price.

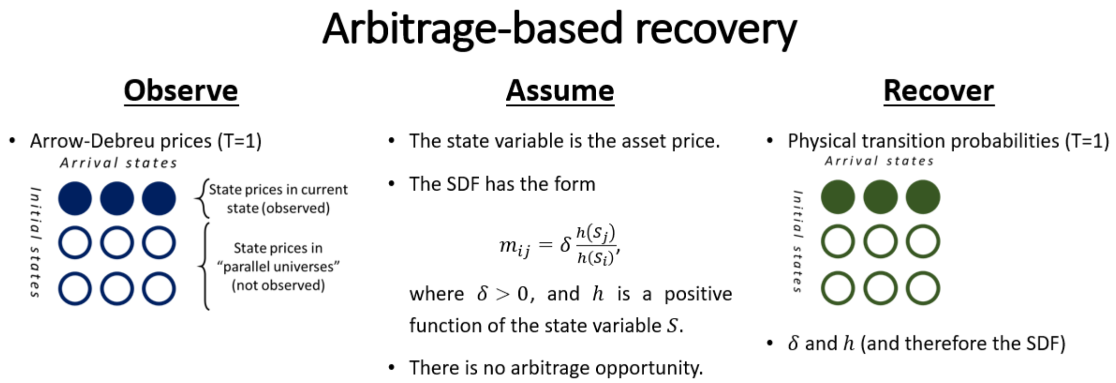

Figure 6: Arbitrage-based recovery. One needs to observe the Arrow-Debreu prices (where the current state is the initial state) for only one maturity. If the SDF takes the form $m_{ij} = \delta \times h(S_j)/h(S_i)$, where $\delta > 0$, S is the asset price and the only state variable, and h is a positive function of S, then in the absence of arbitrage opportunities, the physical transition probabilities (where the current state is the initial state) and the SDF can be recovered.

any other financial asset, such as stock indices, bonds, or currencies. The only state variable in our model is the post-dividend stock price, which can take n different values, collected in the $n \times 1$ column vector $\boldsymbol{S}$. The $n \times n$ matrix of (one-period) Arrow-Debreu prices is $\boldsymbol{A}$, the first row of which corresponds to the current state as the initial state. The stock pays a dividend in one period, which is allowed to be stochastic.[4] We express this dividend as a gross ratio of the post-dividend stock price in the corresponding future state, and collect these dividend ratios in the $n \times n$ matrix $\boldsymbol{F}$. I.e.,

$$F_{ij} = \frac{D_{ij} + S_j}{S_j}, \tag{4}$$

where D_{ij} is the dollar amount of dividend paid in one period if we go from state i to state j.

Before we develop our main theorem, we make a technical assumption about the gross dividend yield matrix $\boldsymbol{F}$, in order to ensure that our recovered SDF takes the form

$$m_{ij} = \delta \frac{h(S_j)}{h(S_i)}, \tag{5}$$

where δ is a positive scalar and h is a positive function of the state variable. Let us

[4]It is also allowed that the stock pays dividends *during* the time period (i.e., not only at the end of the period), as long as those dividends are risk free. We will consider this case in Section 4, when we empirically test our recovery theory observing options on the S&P 500 index.

introduce the scaled gross dividend yield matrix

$$\hat{\boldsymbol{F}} \triangleq \frac{1}{F_{11}} \boldsymbol{F}, \tag{6}$$

recalling that the first row corresponds to the current state as initial state. Hence, F_{11} is the gross dividend yield if the initial state is the current state and the post-dividend stock price remains the same as the current stock price.

Assumption 1. *The scaled gross dividend yield matrix* $\hat{\boldsymbol{F}}$ *is separable in the sense that there exist* $n \times 1$ *column vectors* $\boldsymbol{u} \in \mathbb{R}^n_+$ *and* $\boldsymbol{v} \in \mathbb{R}^n_+$ *such that*

$$\hat{\boldsymbol{F}}' = \boldsymbol{u}\boldsymbol{v}', \tag{7}$$

where

$$v_i = \frac{1}{u_i} \qquad \forall i \in \{1, 2, ..., n\}. \tag{8}$$

This assumption is entirely innocuous regarding the first row of $\hat{\boldsymbol{F}}$ (and thus the first row of $\boldsymbol{F}$),[5] and it restricts only its other rows. But those other rows correspond to "parallel universes", and as such, they are of hardly any concern to us. While developing our arbitrage-based recovery theory, at first we will work with a general positive gross dividend yield matrix $\boldsymbol{F}$, and we impose Assumption 1 only later, when developing Theorem 2.

The first building block of our recovery theorem establishes the uniqueness of the physical transition probabilities and the stochastic discount factor. We formalize this in the following theorem.

Theorem 1. *Given a positive Arrow-Debreu square matrix* $\boldsymbol{A}$ *and a positive gross dividend yield square matrix* $\boldsymbol{F}$ *in a single-period finite-state complete market, there exists a unique decomposition of the Hadamard product* $\boldsymbol{A} \odot \boldsymbol{F}$ *such that* $\boldsymbol{A} \odot \boldsymbol{F} = \phi \boldsymbol{Z} \boldsymbol{\Pi} \boldsymbol{Z}^{-1}$, *where* $\boldsymbol{Z}$ *is a diagonal matrix with positive diagonal elements,* $\boldsymbol{\Pi}$ *is a stochastic matrix,* ϕ *is a positive scalar, and the uniqueness of* $\boldsymbol{Z}$ *is to be understood as "unique up to a positive scale factor".*

Proof. The proof follows the logic of the proof of Result 1 in Martin and Ross (2019). First, we prove the existence of $\boldsymbol{Z}$, $\boldsymbol{\Pi}$, and ϕ. According to the Perron-Frobenius

[5] To see that we can implement any scaled gross dividend yield values in the first row of $\hat{\boldsymbol{F}}$ without violating Assumption 1, let us consider the following. If we want to implement the $1 \times n$ scaled gross dividend yield vector $\hat{\boldsymbol{g}}$ as the first row of $\hat{\boldsymbol{F}}$, then simply define $\boldsymbol{u} \triangleq \hat{\boldsymbol{g}}'$ and the elements of an $n \times 1$ vector $\boldsymbol{v}$ as $v_i = 1/u_i$. Keep in mind that the first element of $\hat{\boldsymbol{g}}$ is, by definition, one, since $\hat{\boldsymbol{g}}$ is a scaled dividend yield vector. Then, the first row of the matrix $\hat{\boldsymbol{F}} = \boldsymbol{v}\boldsymbol{u}'$ is exactly $\hat{\boldsymbol{g}}$, just as we intended, and Assumption 1 holds.

theorem, there exist a positive vector $\boldsymbol{z}$ and a positive scalar ϕ such that $(\boldsymbol{A} \odot \boldsymbol{F})\,\boldsymbol{z} = \phi\boldsymbol{z}$. Define $\boldsymbol{\Pi} \triangleq \frac{1}{\phi}\boldsymbol{Z}^{-1}(\boldsymbol{A} \odot \boldsymbol{F})\,\boldsymbol{Z}$, where $\boldsymbol{Z}$ is a diagonal matrix with the elements of $\boldsymbol{z}$ in its diagonal. Since $\boldsymbol{\Pi}$ is positive and its elements in each of its rows sum to one, it is a stochastic matrix. As one can check, $\boldsymbol{A} \odot \boldsymbol{F} = \phi\boldsymbol{Z}\boldsymbol{\Pi}\boldsymbol{Z}^{-1}$ holds. Thus, there indeed exist ϕ, $\boldsymbol{Z}$, and $\boldsymbol{\Pi}$ satisfying the appropriate conditions in the theorem.

Now, we prove the uniqueness of $\boldsymbol{Z}$, $\boldsymbol{\Pi}$, and ϕ. Multiplying both sides of $\boldsymbol{A} \odot \boldsymbol{F} = \phi\boldsymbol{Z}\boldsymbol{\Pi}\boldsymbol{Z}^{-1}$ by $\boldsymbol{Z}\boldsymbol{\iota}$ from the right (where $\boldsymbol{\iota}$ is the column vector of ones), we obtain $(\boldsymbol{A} \odot \boldsymbol{F})\,\boldsymbol{z} = \phi\boldsymbol{z}$. According to the Perron-Frobenius theorem, ϕ is unique and $\boldsymbol{z}$ is unique up to a positive scale factor. Hence, $\boldsymbol{Z}$ is also unique up to a positive scale factor. And if we rearrange $\boldsymbol{A} \odot \boldsymbol{F} = \phi\boldsymbol{Z}\boldsymbol{\Pi}\boldsymbol{Z}^{-1}$ as $\boldsymbol{\Pi} = \frac{1}{\phi}\boldsymbol{Z}^{-1}(\boldsymbol{A} \odot \boldsymbol{F})\,\boldsymbol{Z}$, we find that therefore $\boldsymbol{\Pi}$ is also unique. ■

Now, we impose our no-arbitrage condition. If there is no arbitrage opportunity in the financial market, then

$$(\boldsymbol{A} \odot \boldsymbol{F})\,\boldsymbol{S} = \boldsymbol{S} \tag{9}$$

must hold. Intuitively, this equation follows from the definition of the Arrow-Debreu matrix $\boldsymbol{A}$ and the gross dividend yield matrix $\boldsymbol{F}$. Namely, if there is no arbitrage opportunity, then the time-zero stock price must be equal to the sum of the Arrow-Debreu prices multiplied by the future stock cash flows (dividend plus post-dividend stock price) in the respective states of the world. I.e.,

$$S_i = \sum_{j=1}^{n} A_{ij} \times \left(F_{ij} \times S_j\right) \qquad \forall\, i \in \{1, 2, \ldots, n\}, \tag{10}$$

which is an equivalent statement to (9). From a mathematical perspective, the no-arbitrage condition (9) explicitly specifies the Perron-Frobenius eigenvector and the corresponding eigenvalue of $\boldsymbol{A} \odot \boldsymbol{F}$. We formalize this in the following lemma.

Lemma 1. *Consider a positive Arrow-Debreu square matrix $\boldsymbol{A}$ and a positive gross dividend yield square matrix $\boldsymbol{F}$ in a single-period finite-state economy. If there is no arbitrage opportunity, then the Perron-Frobenius eigenvector of $\boldsymbol{A} \odot \boldsymbol{F}$ is $\boldsymbol{S}$ and its Perron-Frobenius eigenvalue is* 1.

Proof. The fact that the Perron-Frobenius eigenvector of $\boldsymbol{A} \odot \boldsymbol{F}$ is $\boldsymbol{S}$ and its Perron-Frobenius eigenvalue is 1 readily follows from equation (9) and the Perron-Frobenius theorem. ■

Now, we are ready to state our main theorem, in which we provide our recovered physical transition probabilities and the stochastic discount factor.

Theorem 2 (Arbitrage-Based Recovery Theorem)**.** *Consider a single-period finite-state economy with a positive Arrow-Debreu square matrix* $\boldsymbol{A}$ *and a positive gross dividend yield square matrix* $\boldsymbol{F}$ *satisfying Assumption 1. If there is no arbitrage opportunity, the stochastic discount factor* m_{ij} *and the physical transition probabilities* π_{ij} *are*

$$m_{ij} = \delta \frac{h(S_i)}{h(S_j)}, \tag{11}$$

$$\pi_{ij} = \frac{A_{ij}}{m_{ij}} \tag{12}$$

with $\delta = 1/F_{11}$ *and* $h(S_i) = S_i u_i$*, for all* $i \in \{1, \ldots, n\}$ *and* $j \in \{1, \ldots, n\}$.

Proof. According to Theorem 1, there exist a unique positive scalar ϕ, a unique diagonal matrix $\boldsymbol{Z}$ with positive diagonal elements, and a unique stochastic matrix $\boldsymbol{\Pi}$ such that $\boldsymbol{A} \odot \boldsymbol{F} = \phi \boldsymbol{Z} \boldsymbol{\Pi} \boldsymbol{Z}^{-1}$. Rearranging this, we obtain $\boldsymbol{\Pi} = \frac{1}{\phi} \boldsymbol{Z}^{-1} (\boldsymbol{A} \odot \boldsymbol{F}) \boldsymbol{Z}$. According to Lemma 1, $\phi = 1$ and $Z_{ii} = S_i$. Substituting these values into the rearranged equation, we find $\pi_{ij} = A_{ij} F_{ij} S_j / S_i$. According to (6) and Assumption 1, $F_{ij} = F_{11} u_j / u_i$. Hence, $\pi_{ij} = A_{ij} F_{11} S_j u_j / (S_i u_i)$. Since the SDF is, by definition, equal to the Arrow-Debreu price divided by the physical transition probability, (11) and (12) follow readily. ■

The result of Theorem 2 offers itself as the next step in the evolutionary path of recovery theorems. From a practical perspective, according to the theorem, if there is no arbitrage opportunity and some technical conditions are satisfied, then we can uniquely recover the SDF and the physical transition probabilities in the state space spanned by the asset price. The recovered SDF takes the functional form which is standard also in other recovery theorems: it is the product of a scalar and a function evaluated in the final state, normalized by the same function evaluated in the initial state. Although one might interpret δ as the "subjective discount factor" and the $h(S_i)$ and $h(S_j)$ values as marginal utilities of a pseudo-representative agent, our model does not assume the existence of a representative agent.

Although the form of (11) makes it clear that our recovered SDF takes the same functional form as the SDF of other recovery theorems, sometimes it is easier to work with a slightly different SDF functional form. We provide this (and a corresponding functional form for the recovered physical transition probabilities) in the following corollary.

Corollary 1. *The recovered SDF and physical transition probabilities of Theorem 2 can be equivalently expressed as*

$$m_{ij} = \frac{S_i}{S_j \times F_{ij}} \tag{13}$$

and

$$\pi_{ij} = A_{ij} \times \frac{S_j \times F_{ij}}{S_i} \tag{14}$$

for all $i \in \{1, ..., n\}$ *and* $j \in \{1, ..., n\}$.

Proof. Substituting into δ, $h(S_i)$, and $h(S_j)$ in (11), the result follows immediately. ■

4 Empirical tests

From a theoretical perspective, our Arbitrage-Based Recovery Theorem has several appealing properties. Not only does it require observing only one row of the Arrow-Debreu matrix for only one single maturity, but it does not even assume time-homogeneity. Furthermore, our model explicitly excludes arbitrage opportunities.

Nonetheless, the proof of the pudding is in the eating, and the defense of a theory resides in its empirical validity. Jackwerth and Menner (2020) systematically test to what extent the different versions of recovery theorems are supported empirically. They test three different recovery theorem versions (Ross Original, Ross Basic, Ross Stable/Generalized Recovery), with and without imposing economic restrictions on the Arrow-Debreu price matrix. Based on more than 30 years of S&P 500 options data and using several different density tests (the Berkowitz test, two forms of the Knüppel test, and the Kolmogorov-Smirnov test), they find strong evidence that the 30-day S&P 500 returns do not follow the distribution predicted by the recovery theorems. These findings are also supported by mean-prediction and variance-prediction tests. We test our Arbitrage-Based Recovery Theorem following an approach similar to Jackwerth and Menner (2020). Our results are summarized in Table 2.

4.1 Data

We collect prices of European options written on the S&P 500 index from January 1, 1996 to November 30, 2020 from OptionMetrics. In each month, there is an option expiring on the third Friday[6] of that month. We observe the bid and ask prices of these options 25 days before their expiry.[7] We discard all options which have zero daily trading volume, zero open interest, or a bid price lower than \$0.50. If there are two option

[6] Until February 15, 2015, these options expired on the Saturday immediately following the third Friday of the month, instead of expiring on Friday.

[7] We observe option prices 25 days before expiry instead of 30 days to make sure our model-implied distributions are independent of each other. Since options expire every third Friday of the month, there are 28 days between each expiry. Accounting for no trade taking place on Saturdays and Sundays, we observe the option price on the upcoming Monday, which is 25 days before the expiry date of the option.

observations which are identical apart from one being AM-settled and the other being PM-settled, we keep only the AM-settled option.[8] For each observed option, we calculate the mid price, and based on the put-call parity and the observed spot yield curve[9] on the respective day, we calculate the option-implied dividend yield of the S&P 500 index. We further discard all options which have a negative implied dividend yield or for which the implied dividend yield cannot be calculated (because there is no corresponding put or call observation), or which violate either the lower or the upper option price bound. Moreover, we only keep out-of-money (OTM) options. For each option, we then calculate the implied volatility using the Black-Scholes option-pricing formula. Eventually, our sample consists of 247 observation days.

Using a slightly modified version of the Fast and Stable Method of Jackwerth (2004),[10] we fit a curve on the "observed" implied volatilities for each observation day, using a fine grid of $\Delta = \$0.10$ along the exercise price dimension. Transforming back each implied volatilities along the grid into a European call option price and then taking the (numeric) second derivative of this price with respect to the exercise price, we obtain the state-price density (i.e., the continuous-time equivalent of Arrow-Debreu prices) for each observation day. For details on why the second derivative of a European call option price with regards to the exercise price is equal to the state-price density, we refer to Ross (1976) and Breeden and Litzenberger (1978).

The state-price density is the product of the physical probability density and the stochastic discount factor, and our ultimate goal is to separately identify these two components, based on our Arbitrage-Based Recovery Theorem. In the previous step, for each initial price level S_i and for each possible final price level S_j we have determined the corresponding Arrow-Debreu price A_{ij}.[11] But to apply equation (14), we still need to determine the gross dividend yield F_{ij}. Assuming that the dividend to be received during the life of the option is known at the time of observation (which is not an unreasonable assumption given that we observe the options 25 days before their expiration), the gross dividend yields can be readily calculated from the option-implied dividend yields. Then, substituting into (14), we find the model-implied physical probability density value for each price level.

[8] We do this since the vast majority of our observed options are AM-settled.

[9] Our spot yield curve data is also from OptionMetrics. To obtain the 25-day spot interest rate, we linearly interpolate between the two closest available maturities (e.g., 21 days and 30 days).

[10] When calculating the numerical derivatives of the fitted volatility curve, Jackwerth and Menner (2020) assume that the volatility corresponding to an exercise price just one step outside the boundary is equal to the implied volatility on the boundary. This assumption effectively pushes the first derivative of the volatility curve towards zero on the boundaries. To allow for larger flexibility, our boundary condition instead assumes that the second derivatives on the boundaries are zero.

[11] We can readily calculate the Arrow-Debreu prices by simply multiplying our recovered state-price densities by Δ.

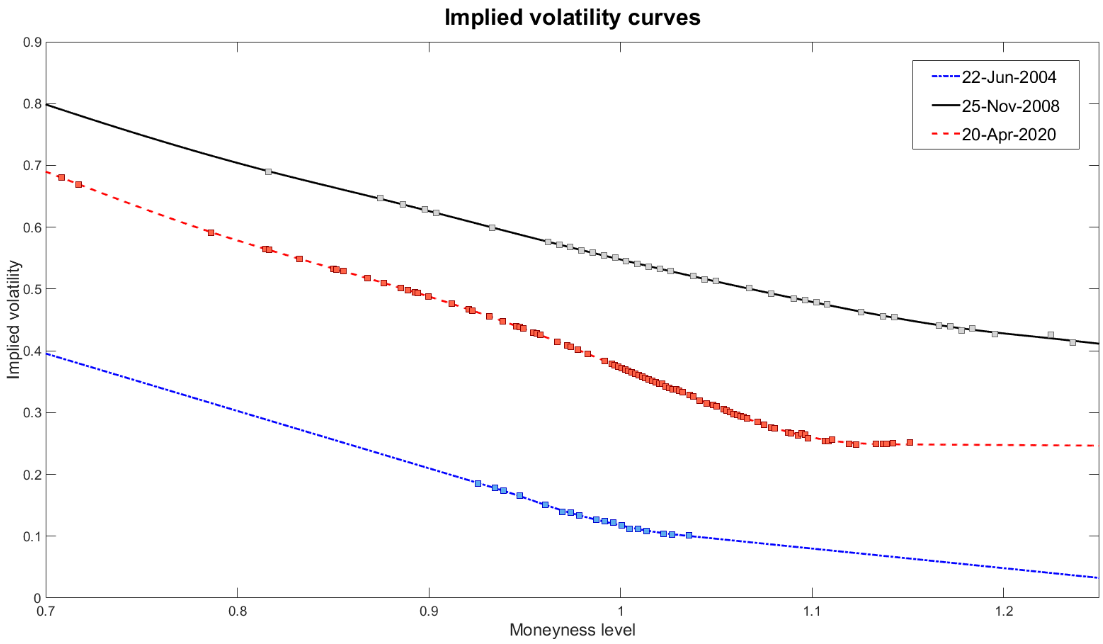

Figure 7: Implied volatility curves on three different dates, corresponding to a 25-day horizon. The markers indicate implied volatilities of observed options.

For demonstration purposes, in Figure 7, we depict our implied volatility curves (together with the "observed" implied volatilities) for three different days: June 22, 2004; November 25, 2008; and April 20, 2020. In Figure 8, we also show the physical probability density functions for these days, recovered by our model. The first date, June 22, 2004, reflects relaxed market conditions, with implied volatilities between 10 and 20 percent for the observed options. The recovered physical probability density function looks "standard": it is centered around a level slightly higher than one, and it features a slight left skew. The two other dates correspond to crisis periods, with much higher model-implied volatilities. By 25 November, 2008, the severity of the global financial crisis became evident, and the strong negative skewness of the density function indicates that the market was prepared for further serious declines in the S&P 500 level. A high (more than 10%) increase in the index level had also much higher probability than on the "relaxed" date of 22 June, 2004, which reflects the huge uncertainty perceived by the market. The situation was similar on 20 April, 2020, by when the seriousness of the COVID-19 pandemic on a global level became apparent. Comparing the two crises dates, the market during the pandemic attributed a much lower probability to a large positive monthly jump in the S&P 500 index level, which is reflected in the difference between the right tails of the two "crisis" distributions.

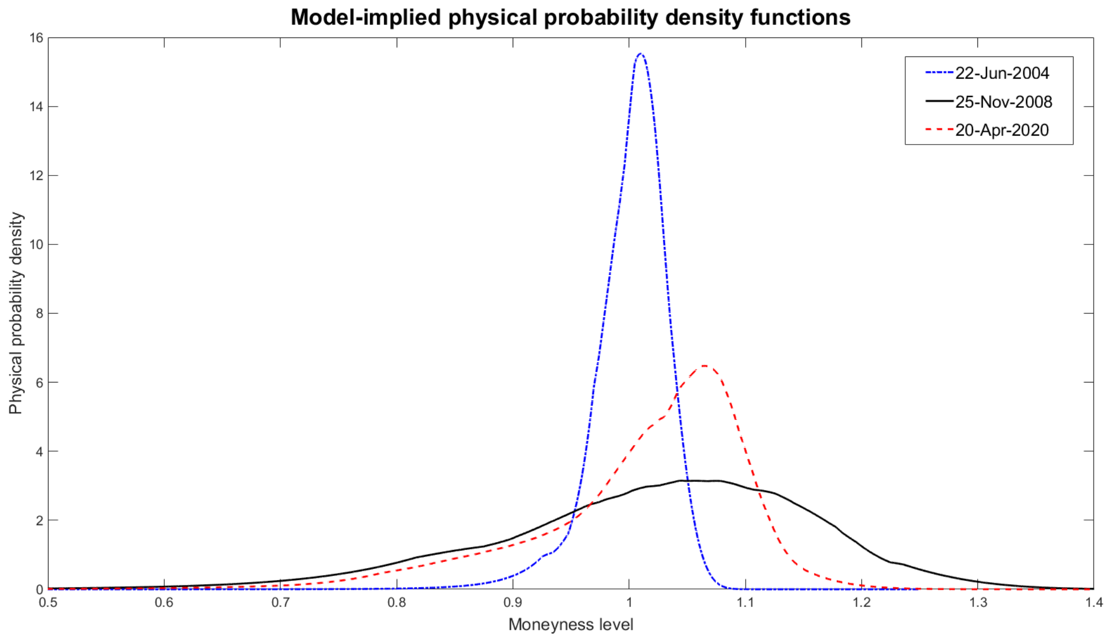

Figure 8: Physical probability density functions on three different dates, corresponding to a 25-day horizon, recovered by our model.

4.2 Density tests

Now, we test whether realized S&P 500 index levels on the option exercise dates are indeed drawn from our model-implied distributions. Following Jackwerth and Menner (2020), we perform several density tests and mean-prediction tests. Our null hypothesis is that the S&P 500 index levels observed on the option exercise dates are realizations of a random variable with our model-implied probability density function.

To perform our density tests, first we observe the S&P 500 index level on each option exercise day.[12] Then, we transform our model-implied probability density functions into cumulative distribution functions (CDFs). For each option, we determine the CDF value corresponding to the realized S&P 500 level. Under our null hypothesis, these CDF values are drawn from a uniform distribution with support $[0, 1]$, and they are independent of each other. To test this hypothesis, we perform the Berkowitz test, two versions of the Knüppel test, and the Kolmogorov-Smirnov test.

[12]When the option is AM-settled with an exercise day on Friday, we observe the opening S&P 500 index level on that Friday. When the option is AM-settled with an exercise day on Saturday, then we observe the closing S&P 500 index level on the Friday just preceding that Saturday. Some observed options are PM-settled with an exercise day on a Friday; for these options, we observe the closing S&P 500 level on that Friday. For a small number of options, the Friday for which we were supposed to observe the S&P 500 level was a public holiday; in such cases, we observed the index level on the Thursday directly preceding that Friday.

4.2.1 Berkowitz test

To perform the test of Berkowitz (2001), we first transform our observed CDF values using the inverse of the cumulative distribution function of the standard normal distribution. I.e., denoting our CDF value observations by x_t, we perform the transformation

$$z_t = \Phi^{-1}(x_t), \tag{15}$$

where Φ denotes the CDF of a standard normally distributed random variable. Then, under the null hypothesis, our observed z_t values follow a standard normal distribution and they are independent of each other. I.e., in the AR(1) model

$$z_t - \mu = \rho\left(z_{t-1} - \mu\right) + \varepsilon_t, \tag{16}$$

under the null hypothesis we have $\mu = 0$, $\rho = 0$, and $Var(\varepsilon_t) = 1$. Using our z_t observations, we obtain the maximum likelihood estimates of μ, ρ, and $Var(\varepsilon_t)$ (denoted by $\hat{\mu}$, $\hat{\rho}$, and $\widehat{Var}(\varepsilon_t)$), and perform a likelihood ratio test where we compare the AR(1) model (16) with parameters $\hat{\mu}$, $\hat{\rho}$, and $\widehat{Var}(\varepsilon_t)$ to the same model with parameters $\mu = 0$, $\rho = 0$, and $Var(\varepsilon_t) = 1$. Performing the Berkowitz test, we find a p-value of 0.1534. Hence, we cannot reject our null hypothesis at the standard significance levels.

4.2.2 Knüppel test

Next, we perform two versions of the test of Knüppel (2015). To this end, we transform our observed x_t values according to

$$y_t = \sqrt{12}\,(x_t - 0.5). \tag{17}$$

Under our null hypothesis, y_t follows a uniform distribution with support $\left[-\sqrt{3}, \sqrt{3}\right]$. We calculate the first three empirical raw moments of our observed y_t values, and form the 3×1 column vector $\boldsymbol{D}_3$ of differences between the empirical and theoretical moments. We also estimate the covariance matrix $\boldsymbol{\Omega}_3$ of $y_t - m_1$, $y_t^2 - m_2$, and $y_t^3 - m_3$, where m_i denotes the i^{th} theoretical raw moment. Under the null hypothesis, the test statistic

$$\alpha_3 = T \times \boldsymbol{D}_3' \boldsymbol{\Omega}_3^{-1} \boldsymbol{D}_3 \tag{18}$$

asymptotically follows a χ^2 distribution with three degrees of freedom. Performing the test, we find a p-value of 0.1334. Repeating the test with the first four moments (instead of the first three), we find a p-value of 0.1041. I.e., based on the Knüppel test, we cannot reject our null hypothesis at any of the standard significance levels.

4.2.3 Kolmogorov-Smirnov test

To perform the Kolmogorov-Smirnov test, we determine the maximum difference between the empirical cumulative distribution function implied by our x_t values and the theoretical cumulative distribution function of the uniform distribution with support $[0,1]$. We find a p-value of 0.0305, which would suggest a rejection of the null hypothesis.

4.3 Mean prediction tests

Besides carrying out density tests, we also test whether our model-implied expected returns indeed predict the realized returns. Concretely, we run the regression

$$r_t = a + b\mu_t + \epsilon_t, \tag{19}$$

where r_t is the realized return at time t (not accounting for dividends), and μ_t is our model-implied expected return (again, without dividends). Under the null hypothesis, the intercept is $a = 0$ and the slope is $b = 1$. Assuming $b = 1$ and testing whether the intercept is statistically significantly different from zero, we find a p-value of 0.1797. Assuming a zero intercept and testing whether the slope is statistically significantly different from one, our p-value is 0.1823. Finally, without any restriction on our regression model, testing the joint hypothesis that $a = 0$ and $b = 1$, we find a p-value of 0.0108. Thus, only the joint test allows for the rejection of our null hypothesis at the standard significance levels, while testing the intercept and the slope separately does not.

4.4 Empirical tests: summary and robustness checks

We summarize the results of our empirical tests in Table 2. Using our full sample, most of our empirical tests do not allow for rejecting the hypothesis that realized returns are drawn from our model-implied probability distributions. The only two tests which suggest rejection of the null hypothesis are the Kolmogorov-Smirnov test (with a p-value of 0.0305) and the joint test of the intercept and the slope in the mean-prediction equation (19) (with a p-value of 0.0108). Considering that using a very similar sample to ours Jackwerth and Menner (2020) find p-values less than 0.05 for most of the other recovery theorems using tests which are identical to ours, our results are promising regarding the empirical plausibility of our Arbitrage-Based Recovery Theorem.

Table 2: Empirical test p-values. We observe prices of European options written on the S&P 500 index from January 1, 1996 to November 30, 2020 with a monthly frequency 25 days before their expiry. After standard cleaning procedures, our full sample consists of 247 observations. We test the null hypothesis that the S&P 500 index levels observed on the option exercise dates are realizations of a random variable with our model-implied probability density function. We perform four different density tests (Berkowitz test, Knüppel test with three and four moments, Kolmogorov-Smirnov test) and three different mean prediction tests on the regression equation (19). We also repeat our empirical tests on subsamples characterized by "volatility smiles" and "volatility skews" (for how we construct these subsamples, see Footnote 13).

	Full sample (247 obs.)	"Volatility smile" dates (122 obs.)	"Volatility skew" dates (125 obs.)
Berkowitz test	0.1534	0.6255	0.0757
Knüppel test (3 moments)	0.1334	0.8262	0.0630
Knüppel test (4 moments)	0.1041	0.9250	0.0148
Kolmogorov-Smirnov test	0.0305	0.7023	0.0021
Intercept test	0.1797	0.7373	0.0993
Slope test	0.1823	0.7424	0.1002
Joint intercept and slope test	0.0108	0.1785	0.0409

To check the robustness of our empirical results, we performed several robustness checks. We changed the lower and the upper boundaries (from their baseline values of 0.1 and 1.5) of the gross return interval on which we recover the physical probabilities; we excluded only the options which have a bid price lower than $0.10 (instead of $0.50); we changed the "fineness" of our grid; and we changed several attributes of our algorithm which fits a smooth volatility curve on the "observed" implied volatilities. We find that our empirical findings are robust to all of these checks.

We also repeat our empirical analysis on a subsample consisting of observation dates characterized by "volatility smiles" (as opposed to by "volatility skews").[13] We

[13] To distinguish volatility smiles from volatility skews, we perform the following exercise. For each observation day, we try to fit a smooth volatility curve on the "observed" implied volatilities using our algorithm. If this fitting exercise can be performed up to a moneyness level of 1.5 without the volatility curve going into the negative area, we consider the observation day to be characterized by a volatility

find the staggering result that our p-values are much higher than in the full sample. Also the only two tests which in the full sample suggested the rejection of our null hypothesis (the Kolmogorov-Smirnov test and the joint intercept and slope test) result in substantially higher p-values now: the p-value of the Kolmogorov-Smirnov test is 0.7023 and the p-value of the joint intercept and slope test is 0.1785.

To confirm that the high p-values of the "volatility smile" dates are not due to the lower number of observations, we repeat our empirical tests on the "volatility skew" dates, the subsample of which has almost the same size as the subsample of the "volatility smile" dates. We indeed find that the p-values of the "volatility skew" dates are much lower than those of both the full sample and the subsample of the "volatility smile" dates.[14] Hence, the fact that the subsample of the "volatility smile" dates has much higher p-values than the full sample is not due to its smaller sample size, and it indeed suggests that our Arbitrage-Based Recovery Theorem works especially well on dates which are characterized by a volatility smile.

5 Conclusion

We develop a novel recovery theorem based on no-arbitrage arguments. Our Arbitrage-Based Recovery Theorem requires the observation of option prices only for one single maturity and only for the current state. Furthermore, it does not assume time homogeneity of either the transition state prices or the transition probabilities, and it precludes arbitrage opportunities.

To test how well our theorem works empirically, we use 25 years of S&P 500 options data with monthly observation frequency, and perform several density tests and mean prediction tests. We find that the vast majority of the tests do not reject that our Arbitrage-Based Recovery Theorem recovers the true distribution of the S&P 500 index level. Furthermore, we find that our theorem works especially well on dates which are characterized by a "volatility smile".

smile. On the other hand, if the volatility curve would extrapolate into the negative area before reaching a moneyness level of 1.5, we consider the observation day to be characterized by a volatility skew.

[14] The only exception is the joint intercept and slope test, which has the lowest p-value in the full sample.

References

F. Audrino, R. Huitema, and M. Ludwig. An Empirical Implementation of the Ross Recovery Theorem as a Prediction Device. *Journal of Financial Econometrics*, (nbz002), 2019. 2, 5, 7

L. Bachelier. Théorie de la spéculation. *Annales scientifiques de l'École Normale Supérieure*, 17:21–86, 1900. 1

G. Bakshi, F. Chabi-Yo, and X. Gao. A Recovery that We Can Trust? Deducing and Testing the Restrictions of the Recovery Theorem. *Review of Financial Studies*, 31 (2):532–555, 2018. 2

R. Bansal and A. Yaron. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004. 2

J. Berkowitz. Testing Density Forecasts, with Applications to Risk Management. *Journal of Business & Economic Statistics*, 19(4):465–474, 2001. 20

F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. 1

J. Borovička, L. P. Hansen, and J. A. Scheinkman. Misspecified Recovery. *The Journal of Finance*, 71(6):2493–2544, 2016. 1

D. T. Breeden and R. H. Litzenberger. Prices of State-Contingent Claims Implicit in Option Prices. *The Journal of Business*, 51(4):621–651, 1978. 17

J. C. Cox, S. A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979. 2

L. P. Hansen and J. A. Scheinkman. Long-Term Risk: An Operator Approach. *Econometrica*, 77(1):177–234, 2009. 1

J. C. Jackwerth. Option-Implied Risk-Neutral Distributions and Risk Aversion. 2004. 17

J. C. Jackwerth and M. Menner. Does the Ross recovery theorem work empirically? *Journal of Financial Economics*, 137(3):723–739, 2020. 2, 3, 4, 5, 8, 9, 16, 17, 19, 21

C. S. Jensen, D. Lando, and L. H. Pedersen. Generalized recovery. *Journal of Financial Economics*, 133(1):154–174, 2019. 1, 2, 4, 5, 9, 10

M. Knüppel. Evaluating the Calibration of Multi-Step-Ahead Density Forecasts Using Raw Moments. *Journal of Business & Economic Statistics*, 33(2):270–281, 2015. 20

I. W. R. Martin and S. Ross. Notes on the yield curve. *Journal of Financial Economics*, 134(3):689–702, 2019. 13

R. Mehra and E. C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985. 2

R. C. Merton. Theory of Rational Option Pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183, 1973. 1

S. Ross. The Recovery Theorem. *The Journal of Finance*, 70(2):615–648, 2015. 1, 2, 4, 5, 6

S. A. Ross. Options and Efficiency*. *The Quarterly Journal of Economics*, 90(1):75–89, 1976. 17