

Consumption and Portfolio Choice with Option-Implied State Prices*

Yacine Aït-Sahalia

Department of Economics
and Bendheim Center for Finance
Princeton University and NBER[†]

Michael W. Brandt

Fuqua School of Business
Duke University and NBER[‡]

This version: November 2007

Abstract

We propose an empirical implementation of the consumption-investment problem using the martingale representation alternative to dynamic programming. Our method is based on the direct observation of state prices from options data. This greatly simplifies the investor's task of specifying the investment opportunity set and inherits the computational convenience of the martingale representation. Our method also makes explicit the economic trade-off between exploiting differences in state prices and probabilities, which generate variation in consumption, and the consumption smoothing induced by risk aversion. Using options-implied information, we find quantitatively different optimal consumption and portfolio policies than those implied by standard return dynamics.

*We thank seminar participants at the CIRANO Conference on Portfolio Choice for their comments and suggestions. Financial support from the NSF under grant SBR-0350772, the Bendheim Center for Finance at Princeton University is gratefully acknowledged.

[†]Princeton, NJ 08544. Phone: (609) 258-4015. E-mail: yacine@princeton.edu.

[‡]Durham, NC 27708. Phone: (919) 660-1948. E-mail: mbrandt@duke.edu.

1. Introduction

Intertemporal consumption and portfolio choice is a daunting problem, requiring as input a complete characterization of the joint distribution of returns across all states of the world from the current date until the end of the investment horizon. Furthermore, professional investment advice is often of limited help because it delivers mostly predictions about mean returns at different horizons. For example, an analyst may give a stock a “near-term hold” or a “long-term buy” recommendation. How can an investor make portfolio and consumption decisions based on such a terse description of the investment opportunity set?

We propose a new empirical approach to address this question. We decompose the portfolio and consumption choice into two separate problems and use different sources of information to get a handle on each. Consider an economy in which the uncertainty is driven by the stochastic movements of a stock and bond index that, in addition to a riskless money market asset, jointly determine the investment opportunity set. At the most abstract level, the investor’s problem consists of choosing how to allocate scarce resources to the different states of the world at all future dates. We use the martingale representation theory of Cox and Huang (1989), Cox and Huang (1991), Karatzas, Lehoczky, and Shreve (1987) and Pliska (1986) to turn this inherently dynamic optimization problem into a static one. This static solution to the portfolio and consumption problem requires two pieces of information. First, the investor has to figure out how expensive one unit of consumption will be in each future state of the world. Second, the investor needs to determine how likely each state is to be realized.

To obtain the first piece of information, the price of a unit of consumption in each future state, we use the market prices of traded options to infer the joint state price density q of stocks and bonds.¹ The resulting option-implied prices of state-contingent consumption bundles allows the investor to determine in which states consumption is relatively cheap or expensive. The investor then allocates consumption to each state in order to maximize expected utility under the budget constraint.

The solution to the static optimization over state-contingent consumption bundles depends

¹Our approach of evaluating the cost of prospective consumption bundle using option-implied information differs from that of Cox and Huang (1989), Cox and Huang (1991) and related theory papers that link the optimal portfolio and consumption choices to the growth-optimal policies under log utility.

on the second piece of information – the likelihood of each state occurring. While everyone in the market is assumed to be a price-taker and hence faces the same state price density q , different investors can have different views about the physical probability distribution of the states, which we denote as p . We therefore present solutions corresponding to a variety of different cases. First, we solve the problem for an investor who expresses beliefs about the Sharpe ratio of stocks and bonds but takes the shape of the physical distribution p , including its second moment, to be the same as that of the state price density q . Second, we consider an investor who assumes a Gaussian shape for the physical density of log returns, corresponding to a standard Geometric Brownian motion benchmark, with volatility matching that of the state price density (i.e., option-implied volatility) and mean calibrated to the investor’s beliefs about the Sharpe ratio of stocks and bonds. Finally, we consider the case in which the shape of p is estimated nonparametrically from historical data.

By construction, the investment opportunity set in our approach is time-varying, as it reflects the variability in the state prices and the probability distributions of the stock and bond index returns at different horizons. The investor’s optimal demand for the risky assets therefore departs from the myopic solution and potentially includes hedging demands. A more original feature of our approach is that it makes explicit the trade-off that is the economic essence of optimal portfolio and consumption choice. On one hand, the investor wants to consume more in states in which the price of consumption is cheap relative to the probability of realizing these states. This effect tends to make the investor’s optimal consumption path respond to changes in the asset prices, because as prices change, so do the state prices and probabilities. On the other hand, deviations across states from a constant consumption path are penalized by the investor’s risk aversion. The higher risk aversion, the more the investor wants to smooth consumption across states, and the less sensitive the optimal consumption becomes to variations in state prices and probabilities, despite the cost of maintaining a constant consumption level across those states.

Option-implied information is naturally suited for the problem at hand because, like the martingale representation approach, it maps out the set of future states of the world into a cross-section of states for which state-contingent consumption can be purchased today. Option markets reveal directly the cost of consumption in each state because, after possibly

some interpolation, we observe market prices for Arrow Debreu securities covering a broad range of future states and maturities.

We are not aware of previous empirical implementations of the martingale representation approach or of the use of option-implied information in the portfolio choice context. There are, however, important examples in the literature of the use of the martingale representation in other theoretical portfolio choice problems. Cvitanić and Karatzas (1995) solve for the optimal portfolio and consumption choice under proportional transaction costs. Wachter (2002) solves for the optimal choice between stocks and cash when the stock returns are predictable by the dividend-to-price ratio in a complete markets setting. Detemple, Garcia, and Rindisbacher (2003) propose a simulation-based approach for dynamic portfolio optimization that is based on the martingale representation.

The remainder of the paper is organized as follows. We start in Section 2 with a description of the theory underlying our approach. In Section 3, we explain how we infer the joint state price density q from the market prices of Standard and Poors (S&P) 500 index and 10-year Treasury bond futures options. In Section 4, we describe how we construct the p density for the different investor's beliefs we consider. In our empirical implementation, described in Section 5, we consider a CRRA investor choosing between consumption and investment in stocks, long-term bonds, and an instantaneously risk-free money market account. As is clear from our theory section though, nothing in our methodology is specific to this particular specification of preferences. Of course, the empirical results would vary with the utility function adopted, often dramatically so.² Section 6 concludes.

2. Portfolio and Consumption Choice

We start with a description of the theoretical problem, focusing on the respective roles played by the state-price density q and the physical density p in the context of an investor's optimal consumption and investment decision. As discussed above, we rely on the martingale representation approach to reduce the dynamic optimization problem to a static one: indi-

²See Aït-Sahalia and Brandt (2001) for the impact of different utility functions on optimal portfolio and consumption choice in a different methodological context.

vidual investors will implement their lifetime consumption and bequest programs through the purchase at time 0 of individual Arrow-Debreu securities. The Arrow-Debreu allocation is identical to that derived using the dynamic optimization method, where investors can trade continuously in frictionless markets.

2.1 Physical and State-Price Densities

Assume that there are $n + 1$ non-redundant assets in the economy; an instantaneously riskless asset with potentially stochastic rate of return r_t and n risky assets whose prices P_t follow an exogenous Markov process. For example, the asset prices could follow

$$dP_t/P_t = \mu_t dt + \sigma_t dZ_t, \quad (2.1)$$

where μ_t and σ_t denote functions of a vector of state variables Y_t and time t , the matrix σ_t has full rank, and Z_t denotes a vector of n independent Brownian motions. But this is only an example, as nothing in the analysis that follows requires continuity of the paths of the asset prices. Any correlation between dP_{it} and dP_{jt} is introduced by the off-diagonal terms in the matrix σ_t . Assuming dynamically complete markets, changes in the state variables driving the uncertainty in the economy can be perfectly hedged using the n assets. We take the price vector as the state variables, so that $Y_t \equiv P_t$.³

Corresponding to the dynamics (2.1), let $p_t(P_t|P_0)$ denote the physical transition density of the state variables (i.e., the conditional density with respect to the Lebesgue measure of the vector P_t at date t given its value P_0 at date 0). The conditional expectation of a stochastic payoff $X_t(P_t)$, which depends on the future realization of P_t , is then given by:

$$E_0[X_t(P_t)] = \int_0^\infty X_t(P_t) p_t(P_t|P_0) dP_t, \quad (2.2)$$

where the integral is n -dimensional.

To rule out arbitrage opportunities among the assets, contingent claims on the assets, and

³There may be more traded assets in the economy than the dimensionality of Z_t , but redundant assets can be perfectly replicated using the n assets and hence do not expand the investment opportunity set.

the money market account, Harrison and Kreps (1979) show that the pricing operator which maps payoffs at date t into prices at date 0 must be linear, continuous, and strictly positive. The Riesz representation theorem characterizes this pricing operator as an expectation with respect to some measure, which we denote as RN . The no-arbitrage cost M_0 of purchasing at date 0 a contingent claim which pays $X_t(P_t)$ at date t is then given by the expected discounted payoff:

$$M_0 = E_0^{RN} \left[\exp \left(- \int_0^t r_\tau d\tau \right) X_t(P_t) \right], \quad (2.3)$$

where the expectation is taken with respect to the so-called risk-neutral measure RN and the payoffs are discounted at the riskfree rate. When the riskfree rate is time-varying but non-stochastic, the discount factor $\exp(-\int_0^t r_\tau d\tau)$ can be pulled outside the expectations, and when the riskfree rate is constant, the discount factor simplifies to $\exp(-rt)$.

However, when the riskfree rate is stochastic, the discount factor inside the expectation makes pricing with the standard risk-neutral measure cumbersome. For that reason, we change this measure to a sequence of new ones, denoted Q_t . Under Q_t , the price of an asset expressed in units of a maturity-matched zero-coupon bond price is a martingale. In contrast, under the more standard risk-neutral measure RN , asset prices are martingales when expressed in units of the money market account.

Let $D_{0,t}$ denote the price at date 0 of a zero-coupon bond with face value \$1 and maturing at date t . Using the measure Q_t , the no-arbitrage cost M_0 is equal to:

$$\frac{M_0}{D_{0,t}} = E_0^{Q_t} \left[\frac{X_t(P_t)}{D_{t,t}} \right]. \quad (2.4)$$

Since $D_{t,t} = 1$, it follows that:

$$M_0 = D_{0,t} E_0^{Q_t} [X_t(P_t)] = D_{0,t} \int_0^\infty X_t(P_t) q_t(P_t|P_0) dP_t, \quad (2.5)$$

where we assume the measure Q_t admits a so-called state-price density (with respect to the Lebesgue measure) denoted as $q_t(P_t|P_0)$. If r_t is non-stochastic, the two measures RN and Q_t are identical, with $D_{0,t} = \exp(-\int_0^t r_\tau d\tau)$. In general, however, the discounting under RN takes place inside the expectation operator, whereas it takes place outside the expectation

operator under Q_t . In exchange for this simplification, we have a sequence of measures Q_t , a different one for each maturity date t , instead of a single measure RN for all dates.

2.2 Assumptions

Inevitably, translating the theoretical martingale representation into an approach that can be implemented in practice requires some simplifying assumptions, which can be viewed as limitations of the present analysis:

- We take the state variables to be the asset prices P_t directly. This is a fairly common assumption in an exchange economy. In our empirical implementation below, we will have two state variables, an equity and a bond index.
- As we will see below, a particular portfolio plays a special role in the martingale representation formulation: this portfolio, with price denoted G_t , is known as the growth optimum portfolio. It is constructed from the assets available to the investor and is such that it maximizes the investor's expected return. We assume that the growth optimal portfolio's price is a function of the asset prices, $G_t = G(P_t, t)$. This is the same assumption as in Theorem 16.1 of Merton (1992), for instance. In general, the function G will be determined as part of an intertemporal general equilibrium solution for the economy, which is outside the scope of this paper. Under this assumption, the growth optimal portfolio is not a separate state variable. Otherwise, this portfolio being an unobservable dynamic trading strategy, any empirical implementation of the martingale representation becomes practically infeasible.
- Traded prices of options provide us with the marginal distributions of the future asset price distributions for the indices. But no derivatives are currently traded with payoffs that link equity and bond returns the way quantos link equity and currency returns, for instance. So, to construct joint distributions for equity and bond indices, we link the options-implied marginal distributions through a copula function. The copula function introduces a correlation parameter between the state variables. We estimate this parameter under the physical distribution – i.e., using the time series of the state variables.

Girsanov's Theorem implies that the correlation parameter is, instantaneously, identical under both the physical and risk neutral distributions. Here, we carry the correlation parameter forward in time. An alternative is to simulate the instantaneous risk neutral dynamics, obtained from the instantaneous estimates, as in Aït-Sahalia, Wang, and Yared (2001), to eliminate the resulting approximation. Comparing the two reveals that the effect of that approximation in the present context is small.

While these assumptions are restrictive, on the other hand, our approach is largely model-free beyond these assumptions. By construction, the state variables are Markovian. But we do not restrict their dynamics further: for instance, they can jump, have a continuous semimartingale part in addition to a jump part, exhibit stochastic volatility, etc.

2.3 The Investor's Problem

Cox and Huang show that in a dynamically complete market, an investor with period t utility function u_t , terminal date T bequest function b_T , and initial wealth W_0 chooses an optimal consumption path $\{C_t; 0 \leq t \leq T\}$ and bequests W_T to maximize:

$$\mathbb{E}_0 \left[\int_0^T u_t(C_t(P_t)) dt + b_T(W_T(P_T)) \right] \quad (2.6)$$

subject to the budget constraint

$$\frac{W_0}{G_0} \geq \int_0^T \frac{D_{0,t}}{G_t} \mathbb{E}_0[C_t(P_t)] dt + \frac{D_{0,T}}{G_T} \mathbb{E}_0[W_T(P_T)] \quad (2.7)$$

and the feasibility constraints that consumption and bequest amounts remain non-negative.

Under the assumption that $G_t = G(P_t, t)$, we have that

$$G_0 p_t(P_t|P_0) = G_t q_t(P_t|P_0) \quad (2.8)$$

as shown in Theorem 16.1 of Merton (1992). Therefore, the investor's problem becomes

$$\begin{aligned} & \mathbb{E}_0 \left[\int_0^T u_t(C_t(P_t)) dt + b_T(W_T(P_T)) \right] \\ &= \int_0^T \int_0^\infty u_t(C_t(P_t)) p_t(P_t|P_0) dP_t dt + \int_0^{+\infty} b_T(W_T(P_T)) p_T(P_T|P_0) dP_T, \end{aligned} \quad (2.9)$$

subject to the budget constraint

$$\begin{aligned} W_0 &\geq \int_0^T D_{0,t} \mathbb{E}_0^{Q_t} [C_t(P_t)] dt + D_{0,T} \mathbb{E}_0^{Q_T} [W_T(P_T)] \\ &= \int_0^T D_{0,t} \int_0^\infty C_t(P_t) q_t(P_t|P_0) dP_t dt + D_{0,T} \int_0^\infty W_T(P_T) q_T(P_T|P_0) dP_T \end{aligned} \quad (2.10)$$

and the feasibility constraints

$$C_t(P_t) \geq 0 \quad \text{for all } 0 \leq t \leq T \text{ and } P_t > 0 \quad (2.11)$$

$$W_T(P_T) \geq 0 \quad \text{for } P_T > 0. \quad (2.12)$$

In words, the investor chooses how much to consume in each possible state P_t at each future date $0 \leq t \leq T$ and how much to bequest in each terminal state P_T , subject to the no-arbitrage cost of the state-contingent consumption path and bequests being less than or equal to the current wealth W_0 . The integrals in the budget constraint reflect the fact that, due to the linearity of the pricing operator, the no-arbitrage cost of any portfolio of contingent claims (including state-contingent consumption and bequest choices) is simply equal to the sum of the costs of the individual components of the portfolio. The sum here is taken across states ($\int_0^\infty \dots dP_t$) and through time ($\int_0^T \dots dt$). The individual costs are evaluated using the separate measures Q_t for each date.

2.4 Optimal Policies

The main appeal of this complete markets formulation is that the optimal state-contingent consumption and bequest policies, denoted $C_t^*(P_t)$ and $W_T^*(P_T)$, do not involve feedback because the dynamics of M_t and P_t are unaffected by the investor's choices. Nevertheless, it is

known from the work of Cox and Huang (1989), Cox and Huang (1991) that the solution is identical to that of the standard Merton (1971) problem where the maximization of the objective (2.9) occurs over consumption $\{C_t; 0 \leq t \leq T\}$ and the portfolio weights $\{\omega_t; 0 \leq t \leq T\}$, subject in the example (2.1) to the wealth dynamics:

$$dW_t = -C_t dt + W_t [(r_t + \omega'_t (\mu_t - r_t)) dt + \omega'_t \sigma_t dZ_t], \quad (2.13)$$

with the constraints $C_t \geq 0$ and $W_t \geq 0$. At date t , the investor consumes C_t and allocates fractions ω_t to the risky asset and the remainder $1 - \sum_{i=1}^n \omega_{it}$ to the riskless asset. In the Merton setting, the dynamic evolution of one of the state variables, the investor's wealth W_t , is endogenously determined. As a result, the solution is recursive and must be solved using dynamic programming.

Because of the absence of feedback in the complete markets formulation, the investor's problem, although dynamic, can be solved as a static optimization using the constrained Lagrangian method of Kuhn and Tucker. With the single multiplier λ_0^B and the continuum of multipliers $\lambda_t^C(P_t)$ and $\lambda_T^W(P_T)$, representing the budget constraint (2.10), consumption nonnegativity constraint (2.11), and wealth nonnegativity constraint (2.12), the Lagrangian for maximizing the expression (2.9) is:

$$\begin{aligned} & \int_0^T \int_0^\infty u_t(C_t(P_t)) p_t(P_t|P_0) dP_t dt + \int_0^\infty b_T(W_T(P_T)) p_T(P_T|P_0) dP_T \\ & + \lambda_0^B \left\{ W_0 - \int_0^T D_{0,t} \int_0^\infty C_t(P_t) q_t(P_t|P_0) dP_t dt - D_{0,T} \int_0^\infty W_T(P_T) q_T(P_T|P_0) dP_T \right\} \\ & + \int_0^T \lambda_t^C(P_t) C_t(P_t) dt + \lambda_T^W(P_T) W_T(P_T). \end{aligned} \quad (2.14)$$

The first-order conditions with respect to the controls C_t and W_T , obtained by setting to zero the partial derivatives of the Lagrangian, are:

$$\frac{\partial u_t(C_t^*(P_t))}{\partial C} p_t(P_t|P_0) + \lambda_t^C(P_t) - \lambda_0^B D_{0,t} q_t(P_t|P_0) = 0 \quad (2.15)$$

$$\frac{\partial b_T(W_T^*(P_T))}{\partial W} p_T(P_T|P_0) + \lambda_T^W(P_T) - \lambda_0^B D_{0,T} q_T(P_T|P_0) = 0 \quad (2.16)$$

and the first-order conditions with respect to the multipliers are:

$$\begin{aligned}\lambda_t^C (P_t) C_t^* (P_t) &= 0 \\ \lambda_T^W (P_T) W_T^* (P_T) &= 0\end{aligned}\tag{2.17}$$

$$\lambda_0^B \left\{ W_0 - \int_0^T D_{0,t} \int_0^\infty C_t^* (P_t) q_t (P_t|P_0) dP_t dt - D_{0,T} \int_0^\infty W_T^* (P_T) q_T (P_T|P_0) dP_T \right\} = 0.$$

Since $\partial u_t / \partial C > 0$ and $\partial b_T / \partial W > 0$, the investor is unsatiated for both consumption and bequests. As a result, $\lambda_0^B > 0$ and the budget constraint (2.10) is binding. Solving the first order conditions (2.15)-(2.16) for the multipliers λ_t^C and λ_t^W yields:

$$\lambda_t^C (P_t) = \max \left[0, \lambda_0^B D_{0,t} q_t (P_t|P_0) - \frac{\partial u_t (0)}{\partial C} p_t (P_t|P_0) \right]\tag{2.18}$$

$$\lambda_T^W (P_T) = \max \left[0, \lambda_0^B D_{0,T} q_T (P_T|P_0) - \frac{\partial b_T (0)}{\partial W} p_T (P_T|P_0) \right],\tag{2.19}$$

where the max operators reflect the fact that the multipliers are non-zero only when the corresponding choice variables are zero.

Since $\partial^2 u_t / \partial C^2 < 0$ and $\partial^2 b_T / \partial W^2 < 0$, the inverse functions $(\partial u_t / \partial C)^{-1}(\partial u_t / \partial C) = C$ and $(\partial b_T / \partial W)^{-1}(\partial b_T / \partial W) = W$ are well-defined for all $C \geq 0$ and $W \geq 0$ and are strictly decreasing. Solving the first order conditions (2.15)-(2.16), given the budget constraint multiplier λ_0^B , yields the optimal policies:

$$C_t^* (P_t) = \max \left[0, \left(\frac{\partial u_t}{\partial C} \right)^{-1} \left(\lambda_0^B D_{0,t} \frac{q_t (P_t|P_0)}{p_t (P_t|P_0)} \right) \right]\tag{2.20}$$

$$W_T^* (P_T) = \max \left[0, \left(\frac{\partial b_T}{\partial W} \right)^{-1} \left(\lambda_0^B D_{0,T} \frac{q_T (P_T|P_0)}{p_T (P_T|P_0)} \right) \right],\tag{2.21}$$

where here the max operators reflect the fact that either $C_t^* (P_t) = 0$ or $C_t^* (P_t) > 0$, but in the latter case $\lambda_t^C (P_t) = 0$, and similarly for $W_T^* (P_T)$ and $\lambda_T^W (P_T)$.

From equations (2.20)-(2.21), the optimal policies are fully characterized once the (scalar) budget constraint multiplier λ_0^B is determined. Since non-satiation implies that the budget

constraint is binding, plugging the optimal policies (2.20)-(2.21) into equation (2.10) yields:

$$W_0 = \int_0^T D_{0,t} \int_0^\infty \max \left[0, \left(\frac{\partial u_t}{\partial C} \right)^{-1} \left(\lambda_0^B D_{0,t} \frac{q_t(P_t|P_0)}{p_t(P_t|P_0)} \right) \right] q_t(P_t|P_0) dP_t dt + D_{0,T} \int_0^\infty \max \left[0, \left(\frac{\partial b_T}{\partial W} \right)^{-1} \left(\lambda_0^B D_{0,T} \frac{q_T(P_T|P_0)}{p_T(P_T|P_0)} \right) \right] q_T(P_T|P_0) dP_T, \quad (2.22)$$

which determines λ_0^B . Replacing λ_0^B by its value in equations (2.20)-(2.21) completes the characterization of the investor's optimal policies.

In the special but popular case of constant relative risk aversion (CRRA) utility with $u_t(C) = \beta^t C^{1-\gamma}/(1-\gamma)$ and $b_T(W) = \beta^T W^{1-\gamma}/(1-\gamma)$, the inverse functions are:

$$\left(\frac{\partial u_t}{\partial C} \right)^{-1}(u) = \left(\frac{\beta^t}{u} \right)^{1/\gamma} \quad \text{and} \quad \left(\frac{\partial b_T}{\partial W} \right)^{-1}(b) = \left(\frac{\beta^T}{b} \right)^{1/\gamma}. \quad (2.23)$$

The budget constraint (2.22) therefore reduces to $W_0 = (\lambda_0^B)^{-1/\gamma} I_0$ with:

$$I_0 \equiv \int_0^T D_{0,t} \int_0^\infty \max \left[0, \left(\frac{\partial u_t}{\partial C} \right)^{-1} \left(D_{0,t} \frac{q_t(P_t|P_0)}{p_t(P_t|P_0)} \right) \right] q_t(P_t|P_0) dP_t dt + D_{0,T} \int_0^\infty \max \left[0, \left(\frac{\partial b_T}{\partial W} \right)^{-1} \left(D_{0,T} \frac{q_T(P_T|P_0)}{p_T(P_T|P_0)} \right) \right] q_T(P_T|P_0) dP_T. \quad (2.24)$$

Moreover, the optimal policies can be written in terms of the consumption to initial wealth and terminal wealth to initial wealth ratios as:

$$\frac{C_t^*(P_t)}{W_0} = \frac{1}{I_0} \max \left[0, \beta^{t/\gamma} \left(D_{0,t} \frac{q_t(P_t|P_0)}{p_t(P_t|P_0)} \right)^{-1/\gamma} \right] \quad (2.25)$$

$$\frac{W_T^*(P_T)}{W_0} = \frac{1}{I_0} \max \left[0, \beta^{T/\gamma} \left(D_{0,T} \frac{q_T(P_T|P_0)}{p_T(P_T|P_0)} \right)^{-1/\gamma} \right]. \quad (2.26)$$

2.5 Relative Price Differences vs. Consumption Smoothing

Equations (2.25)-(2.26), or their more general versions (2.20)-(2.21), illustrate the trade-off that is the economic essence of optimal portfolio and consumption choice. On one hand, the investor wants to consume more in states in which the price of consumption is cheap relative

to the probability of realizing these states (choose a higher C_t^* when the ratio q_t/p_t is low). This effect tends to make the investor's optimal consumption path respond to changes in the asset prices, because as P_t changes, so does q_t/p_t . On the other hand, deviations across states from a constant value $C_t^* = C^*$ are penalized by the investor's risk aversion. The higher γ , the more the investor wants to smooth consumption across states and the less sensitive C_t^* becomes to variations in q_t/p_t . In the limit as $\gamma \rightarrow \infty$, the optimal policy becomes $C_t^* = C^*$, irrespectively of the price of consumption in different states.

2.6 Portfolio Implementation of the Optimal Investment Policy

Once the optimal state contingent consumption and bequest plan has been determined, this plan can be implemented at date 0 by purchasing pure Arrow-Debreu securities. Specifically, to implement the optimal consumption plan, the investor purchases for every state P at each future date t a quantity $C_t^*(P) dt$ of Arrow-Debreu securities paying \$1 if $P_t = P$ and 0 otherwise. In addition, the investor buys quantities $W_T^*(P)$ of Arrow-Debreu securities for states P at the terminal date T , to implement the optimal bequest plan.

Arrow-Debreu securities can be synthesized or at least approximated closely by butterfly strategies involving plain vanilla European call options on the underlying assets. Consider the following butterfly strategy payoff:

$$X_t(P_t; K, \varepsilon) = \frac{\max[0, P_t - K + \varepsilon] + \max[0, P_t - K - \varepsilon] - 2 \max[0, P_t - K]}{\varepsilon^2}, \quad (2.27)$$

formed using European call options with strike prices $K - \varepsilon < K < K + \varepsilon$. A security with this payoff converges to an Arrow-Debreu security at $P = K$ in the limit as $\varepsilon \rightarrow 0$. It follows immediately that the optimal consumption and bequest plan can be implemented by trading a basket of European call options on the underlying assets.

In case the required call options are not directly available in a liquid market, they can themselves be replicated by a dynamic trading strategy in the underlying assets. It is precisely this replicating strategy in the underlying assets which most of the portfolio choice literature focuses on. The economic point of the preceding discussion is that, in our complete markets framework, the optimal portfolio choice is fully characterized by the optimal state contingent

consumption and bequest plan. As a result, the remainder of the paper focuses on the latter economic decision.

3. Option-Implied State Price Density q

The previous section showed that the optimal state contingent consumption and bequest plan as well as its trading implementation are fully determined by functions of the state-price densities $q_t(P_t|P_0)$ and the physical transition densities $p_t(P_t|P_0)$. We now discuss how to characterize empirically these two objects, starting with the state-price densities q_t . In a nutshell, we use data on exchange traded European put and call options to infer the state-price densities using a parametric multivariate counterpart to the nonparametric univariate method described in Aït-Sahalia and Lo (1998).

3.1 Marginal SPD for Each Asset Class

Consider first the case of a single risky asset and assume initially that the asset price P_t is distributed log-normal under the measure Q_t with a mean of $E_0^{Q_t}[P_t]$ and a volatility of log returns $\ln(P_t) - \ln(P_0)$ equal to $\sigma_{0,t}\sqrt{t}$. We refer to this model as the Black-Scholes case, although in the standard Black-Scholes model P_t is log-normal under the single risk-neutral measure RN , as opposed to the sequence of measures Q_t .⁴ Define the yield of a zero-coupon bond with maturity date t to be $Y_{0,t}$, so that $D_{0,t} = \exp(-Y_{0,t}t)$.

Let H denote the price of a European call option with maturity date t and strike price K , given by equation (2.5) evaluated for the payoff function $X_t(P_t) = \max(P_t - K, 0)$. Under the Black-Scholes assumptions, we have:

$$\begin{aligned} H_{BS}(P_0, K, t, \sigma_{0,t}) &= D_{0,t} E_0^{Q_t} [\max(P_t - K, 0)] \\ &= D_{0,t} (F_{0,t}\Phi(d_1) - K\Phi(d_2)), \end{aligned} \tag{3.1}$$

⁴Under the Black-Scholes assumption of a constant interest rate, $r_t = r$, which we do not adopt here, the two sets of measures are identical.

with

$$d_1 \equiv \frac{\ln(F_{0,t}/K)}{\sigma_{0,t}\sqrt{t}} + \frac{1}{2}\sigma_{0,t}\sqrt{t} \quad \text{and} \quad d_2 \equiv d_1 - \sigma_{0,t}\sqrt{t}. \quad (3.2)$$

$F_{0,t}$ denotes the price of a forward contract for delivery of the asset at time t , which, using equation (2.5), equals the expected future spot price under the measure Q_t . If the asset pays income at a rate of $\delta_{0,t}$, this forward price is $F_{0,t} = E_0^{Q_t}[P_t] = P_0 \exp((Y_{0,t} - \delta_{0,t})t)$.

The state price density q_t in the Black-Scholes case is given by:

$$q_{BS,t}(P_t|P_0) = \frac{1}{P_t\sqrt{2\pi t}\sigma_{0,t}} \exp\left(-\frac{(\ln(P_t/P_0) - (Y_{0,t} - \delta_{0,t} - \sigma_{0,t}^2/2)t)^2}{2t\sigma_{0,t}^2}\right) \quad (3.3)$$

and the corresponding physical transition density p_t is:

$$p_{BS,t}(P_t|P_0) = \frac{1}{P_t\sqrt{2\pi t}\sigma_{0,t}} \exp\left(-\frac{(\ln(P_t/P_0) - (\mu_{0,t} - \sigma_{0,t}^2/2)t)^2}{2t\sigma_{0,t}^2}\right), \quad (3.4)$$

where $\mu_{0,t}$ denotes the expected rate of return on the asset between times 0 and t . This expected rate of return is defined indirectly by the equation $E_0[P_t] = P_0 \exp(\mu_{0,t}t)$.

Suppose now, as is the overwhelmingly common assumption in practice, that the call pricing function is given by the Black-Scholes formula (3.1), except that the volatility parameter for a given option is determined by a function $\sigma_{0,t} = \sigma_0(K/F_{0,t}, t)$ of the moneyness $M_{0,t} \equiv K/F_{0,t}$ and time-to-maturity t of the option:

$$H(P_0, K, t) = H_{BS}(P_0, K, t, \sigma_0(K/F_{0,t}, t)). \quad (3.5)$$

Applying the basic pricing equation (2.5) to this far more general case yields:

$$H(P_0, K, t) = D_{0,t} E_0^{Q_t}[\max(P_t - K, 0)] = D_{0,t} \int_K^{+\infty} (P_t - K) q_t(P_t|P_0) dP_t. \quad (3.6)$$

Following Banz and Miller (1978) and Breeden and Litzenberger (1978), the state price density can then be recovered by direct differentiation:

$$q_t(K|P_0) = \frac{1}{D_{0,t}} \frac{\partial^2 H}{\partial K^2}(P_0, K, t), \quad (3.7)$$

where the total derivatives account for the dependence of the function $\sigma_0(K/F_{0,t}, t)$ on K :

$$\begin{aligned}\frac{\partial H}{\partial K} &= \frac{\partial H_{BS}}{\partial K} + \frac{1}{F} \frac{\partial \sigma_0}{\partial M} \frac{\partial H_{BS}}{\partial \sigma} \\ \frac{\partial^2 H}{\partial K^2} &= \frac{\partial^2 H_{BS}}{\partial K^2} + \frac{2}{F} \frac{\partial \sigma_0}{\partial M} \frac{\partial^2 H_{BS}}{\partial K \partial \sigma} + \frac{1}{F^2} \left(\frac{\partial \sigma_0}{\partial M} \right)^2 \frac{\partial^2 H_{BS}}{\partial \sigma^2} + \frac{1}{F^2} \frac{\partial^2 \sigma}{\partial M^2} \frac{\partial H_{BS}}{\partial \sigma}.\end{aligned}\tag{3.8}$$

Aït-Sahalia and Lo (1998) exploit this equation to infer an estimate of q_t from a non-parametric estimate of the second partial derivative of the call pricing function with respect to the strike price. We follow a similar approach, except that we parametrically model the volatility function $\sigma_0(K/F_{0,t}, t)$. We discuss our specific modelling choice in the context of the empirical application.

3.2 Joint SPD for Two Asset Classes

We now extend this idea to two assets, say stocks and bonds, with price vector $P_t = (P_{1t}, P_{2t})'$. The price of a call option written on the first asset is, again from equation (2.5):

$$\begin{aligned}H_1(P_0, K, t) &= D_{0,t} \int_0^{+\infty} \int_K^{+\infty} (P_{1t} - K) q_t(P_t|P_0) dP_{1t} dP_{2t} \\ &= D_{0,t} \int_K^{+\infty} (P_{1t} - K) q_{1t}(P_{1t}|P_0) dP_{1t},\end{aligned}\tag{3.9}$$

where we denote by $q_{1t}(P_{1t}|P_0)$ the marginal state price density:

$$q_{1t}(P_{1t}|P_0) = \int_0^{+\infty} q_t(P_{1t}, P_{2t}|P_0) dP_{2t}.\tag{3.10}$$

Applying the univariate procedure described in the previous section to options on the first asset alone therefore allows us to estimate $q_{1t}(P_{1t}|P_0)$. Similarly, the marginal state price density $q_{2t}(P_{2t}|P_0)$ can be estimated from options on the second asset alone.

To obtain a joint distribution from the two marginal densities, it is convenient to first transform variables from prices to log returns. Let R_t denote the annualized log return implied by the prices P_t . The two are related by the deterministic change of variable $P_t = P_0 \exp(t R_t)$,

so that the density of returns is given by the Jacobian formula:

$$q_t(R_t|P_0) = t P_0 \exp(R_t) q_t(P_t|P_0). \quad (3.11)$$

Since there is no risk of confusing the densities of prices and returns, we do not distinguish the notation between the two. The argument (P_t or R_t) indicates which is which.

We assume that the joint state price density $q_t(R_t|P_0)$ implied by the two marginal state price densities is of the Plackett (1965) form.⁵ Specifically, we assume that the joint cumulative distribution function (CDF):

$$Q_t(R_t|P_0) = \int^{R_t} q_t(R_t|P_0) dR_t \quad (3.12)$$

is given by:

$$Q_t(R_t|P_0) = \frac{\{A_t(R_t|P_0)^2 + 4\theta(1-\theta)Q_{1t}(R_{1t}|P_0)Q_{2t}(R_{2t}|P_0)\}^{1/2} - A_t(R_t|P_0)}{2(1-\theta)}, \quad (3.13)$$

where Q_{1t} and Q_{2t} are the CDFs corresponding to the marginal densities q_{1t} and q_{2t} , and:

$$A_t(R_t|P_0) = 1 - (1-\theta)\{Q_{1t}(R_{1t}|P_0) + Q_{2t}(R_{2t}|P_0)\}. \quad (3.14)$$

The Plackett family is parameterized by θ , which controls the correlation between the two variables R_{1t} and R_{2t} . In particular, $\theta = 0$ corresponds to a correlation $\rho = -1$, $\theta = 1$ to a correlation $\rho = 0$, and $\lim_{\theta \rightarrow +\infty}$ to a correlation $\rho = 1$. In the special case of $\theta = 1$, which yields uncorrelated variables, equation (3.13) turns into:

$$Q_t(R_t|P_0) = Q_{1t}(R_{1t}|P_0)Q_{2t}(R_{2t}|P_0). \quad (3.15)$$

In our empirical application, we calibrate the parameter θ to the historical correlation of log returns over horizon t .⁶ Finally, given the joint CDF (3.12) from the Plackett formula, we

⁵See Rosenberg (2003) for another use of the Plackett family of densities with given marginal densities.

⁶By Girsanov's Theorem, the second moments are unaffected by the change of measure in the continuous-time limit, and approximately so at finite horizons. This justifies using the empirical correlation as a proxy

recover the corresponding joint density using:

$$q_t(R_t|P_0) = \frac{\partial^2 Q_t}{\partial R_{1t} \partial R_{2t}}(R_t|P_0). \quad (3.16)$$

4. Physical State Density p

While the state price densities $q_t(P_t, P_0)$ can be inferred objectively from the prices of traded options, the corresponding physical densities $p_t(P_t, P_0)$ are by nature dependent on the subjective views of a particular investor. We discuss here three different ways of constructing p_t , which differ primarily in the weight placed on historical market data, current market data, and subjective beliefs.

4.1 Belief-Induced Shifts of the SPD

We first consider an investor who, in light of Girsanov's Theorem, takes the shapes of the physical densities p_t to be the same as the shapes of the corresponding state price densities q_t , inferred from option prices as described above, and only expresses a view about the Sharpe ratios or risk premia (given second moments from the state price densities – recall the argument in footnote 6) of the two assets. Specifically, given a vector of annualized risk premia μ_t for horizon t , we set:

$$p_t(R_t|P_0) = q_t(R_t - \mu_t t|P_0). \quad (4.1)$$

Since the option-implied state price density is forward-looking and conditional on current information, this specification of p_t incorporates heteroscedasticity and time-variation in higher-order moments. It does so, however, using the information contained in option prices, as opposed to requiring the investor to build a sophisticated statistical model for returns. This simplification is one of the key advantages of our general approach.

The investor can shift the state-price density either by historical risk premia estimates or impose a subjective belief about future expected returns. This subjective belief could

for the risk-neutral correlation. See Aït-Sahalia, Wang, and Yared (2001) for a different use of this argument to construct a test of the hypothesis that the state price density q_t accurately prices cross-sectional options given the time series evidence on the underlying asset.

be formed through a forecasting model, professional investment advice, introspection, or a combination thereof.⁷ Notice that in the limiting case where the investor believes that the risk premia are zero (so $p_t = q_t$), the optimal policies are state-independent, since equations (2.20)-(2.21) reduce to:

$$C_t^*(P_t) = \max \left[0, \left(\frac{\partial u_t}{\partial C} \right)^{-1} (\lambda_0^B D_{0,t}) \right] \quad (4.2)$$

$$W_T^*(P_T) = \max \left[0, \left(\frac{\partial b_T}{\partial W} \right)^{-1} (\lambda_0^B D_{0,T}) \right]. \quad (4.3)$$

4.2 Gaussian Density

We also consider the case of Gaussian physical densities:

$$p_t(R_t|P_0) = \mathcal{N}[(Y_{0,t} + \mu_t)t, \Sigma_t t] \quad (4.4)$$

where the vector of annualized risk premia μ_t and the annualized return covariance matrix Σ_t for horizon t are both specified by the investor, and where $\mathcal{N}[a, b]$ denotes the Gaussian density with mean a and variance b . One possible justification for this relatively simple specification of p_t is that a typical investor, already finding it challenging to form views about the first two moments of returns, is unlikely to be willing or able to express strong beliefs about the skewness, kurtosis, and higher order moments of the (log) return density. The investor therefore defaults to the intuitive and theoretically appealing Black-Scholes benchmark case discussed above.

As in the case of belief-induced shifts of the state price density, the moments μ_t and Σ_t of the Gaussian density can be based on a forecasting model, professional investment advice, introspection, or a combination thereof. Alternatively, the covariance matrix Σ_t can be calibrated, again in light of Girsanov's Theorem, to the second moments of the state price density, which leaves only the risk premia to be specified by investor. In the tables and figures below, we will express views directly on the Sharpe ratios of the two assets.

⁷For instance, risk premia can be calibrated to the consensus forecasts by the academic finance profession or by Chief Financial Officers, as reported in Welch (2000) and Graham and Harvey (2002), respectively.

4.3 Empirical Density

Finally, an obvious case can be made for using historical data as the basis for constructing the physical density p_t . Given a time series of realized log returns over horizon t , $R_{1,t,s}$ and $R_{2,t,s}$, for $s = 1, \dots, S$, we construct for each asset a histogram of the marginal physical distribution $p_{it}(P_{it}|P_{i0})$ or a smoothed version of it in the form of a kernel density estimator.⁸ We then obtain the joint physical density $p_t(P_t|P_0)$ by combining the two marginal densities using the Plackett formula described in the previous section. The parameter θ is again calibrated to the historical correlation of log returns at horizon t .

In order to focus on the shape of the physical densities, we rescale in our application the empirical densities to have the same second moments as the state-price densities and shift them according to views on the Sharpe ratios of the two assets. The only difference between using empirical densities and the other two cases discussed above is that the shape of the empirical densities is by construction backward-looking and unconditional. In contrast, the shape of the state-price densities is forward-looking and conditional on current information.

5. Empirical Implementation

We now implement our approach empirically. We use option prices to infer the state price densities q_t , and then examine the optimal consumption policies corresponding to the different choices of the physical state densities p_t described above. Besides illustrating the mechanics of our approach, the contribution of this application is two-fold. First, we investigate empirically the trade-off between exploiting differences in the prices of consumption in different states and smoothing consumption across states, for different choices of the physical state density. Second, we illustrate the tension between the state prices inferred from option prices and the beliefs expressed in the risk premium surveys mentioned above. One can view this tension as another take on the equity premium puzzle.

We consider an individual who can invest in two risky securities, a stock fund that tracks the S&P 500 index and a bond fund with duration equal to that of a 10-year Treasury note

⁸See, for example, Wand and Jones (1995) for a description of kernel density estimation techniques.

futures, in addition to the risk-free money market account. We solve for the optimal consumption policies assuming the investor has CRRA preferences with $\gamma = 5$ and $\beta = 0.98$. For simplicity, we abstract from the bequest motive by assuming $b_T(W_T) = 0$.

Although we focus on CRRA preferences in this application, it is clear from the theory that our approach can be applied just as easily to any other choice of preferences. Given estimates of q_t and p_t , the optimal consumption choice is characterized by equation (2.20), which simplifies to equation (2.25) in the case of CRRA preferences. The estimation of the densities is independent of the investor's preferences. It follows that the relative prices of consumption in different states, which depend only on the ratio q_t/p_t , are also unaffected by the preferences. The only role played by the utility function (its inverse actually) is the relative desire of the investor to smooth consumption across dates and states.

5.1 Data

We collect options data for the S&P 500 index (SPX) and the 10-year Treasury note futures (TY). The SPX options are traded at the Chicago Board Options Exchange (CBOE) and the TY options are traded at the Chicago Board of Trade (CBOT). Both are extremely liquid, with aggregate daily volumes well in excess of 100,000 contracts each. For 10 consecutive days (January 6–17, 2003) we obtain the bid and ask quotes of all listed SPX and TY options at 11:30 am. We then take the price of each option to be the midpoint between the bid and ask quotes. Finally, we compute implied volatilities from these mid-point prices using the prevailing term structure of Eurodollar interest rates. Table 1 describes the options data.

We also collect daily closing prices of the S&P 500 index and the on-the-run 10-year Treasury note from January 2, 1962 through January 17, 2003. We use this data to compute the historical moments and the joint empirical densities of the stock and bond returns at different horizons. Table 2 describes the underlying asset data.

5.2 Density Estimation

We use the raw options data to infer first the marginal and ultimately the joint state price densities for the S&P 500 index and 10-year Treasury bond. The SPX option contract

is European style, as required by our econometric approach. The TY option contract, in contrast, is American style, allowing for early exercise prior to expiration. The underlying asset of the TY contrast is a 10-year Treasury note futures contract; see Fama and French (1988) for an adjustment for the early exercise premium, and the construction of an equivalent European option.

In-the-money options are notoriously illiquid and their prices can be unreliable. We therefore replace the in-the-money option prices with ones implied by put-call-parity and more liquid out-of-the-money options. Given a European call option price $H(P_{i0}, K, t)$, put-call-parity yields the European put option price $L(P_{i0}, K, t)$ for the same strike price and maturity date:

$$H(P_{i0}, K, t) - L(P_{i0}, K, t) = D_{0,t} \left(\mathbb{E}_0^{Q_t} [P_{it}] - K \right). \quad (5.1)$$

We first evaluate put-call parity at the money ($K \simeq P_{i0}$), where both the call and put are the most liquid, to obtain a common implied forward price $F_{i0} = \mathbb{E}_0^{Q_t} [P_{it}]$.⁹ We then apply this option-implied forward price to put-call-parity for all away-from-the-money strike prices to replace the illiquid in-the-money option prices with liquid out-of-the-money ones.

We next invert each option price to an implied volatility, using the at-the-money option-implied forward price of the underlying asset, and fit a standard parametric implied volatility surface (the so-called "practitioner Black-Scholes model") by constrained ordinary least squares. Specifically, we model the implied volatilities for the S&P 500 index as:

$$\begin{aligned} \sigma_0(K/F_{10,t}, t) &= 0.696 - 0.308 (K/F_{10,t}) - 0.175 t + 0.132 (K/F_{10,t}) t \quad \text{if } t \leq 2.5 \\ &= 0.259 \quad (\text{the fitted value for } t = 2.5) \quad \text{if } t > 2.5, \end{aligned} \quad (5.2)$$

and the implied volatilities for the 10-year Treasury note futures as:

$$\begin{aligned} \sigma_0(K/F_{20,t}, t) &= 0.356 - 0.265 (K/F_{20,t}) - 0.158 t + 0.265 (K/F_{20,t}) t \quad \text{if } t \leq 1.0 \\ &= 0.197 \quad (\text{the fitted value for } t = 1.0) \quad \text{if } t > 1.0. \end{aligned} \quad (5.3)$$

The implied volatility surface is constrained to be flat, with continuous first and second deriv-

⁹The forward price is the value $F_{i0,t}$ that solves $D_{0,t} \mathbb{E}_0^{Q_t} [P_{it} - F_{i0,t}] = 0$.

atives, at a horizon of 2.5 years for the S&P 500 index and one year for the 10-year Treasury note futures. The reason for these constraints is that liquid options are not available beyond these maturities. The constraints effectively set the implied volatilities for longer horizons equal to the implied volatilities of the longest available at-the-money options. The model specifications fit the data well, with multiple R^2 equal to 95.5% for the S&P 500 index and 79.5% for the 10-year Treasury note futures options.

The advantage of these simple parametric models is that the partial derivatives needed for equation (3.8) can be evaluated analytically.¹⁰ Figure 1 plots the resulting marginal SPDs q_t for horizons t ranging from one to 10 quarters (2.5 years). The figure also shows deviations of the SPDs from Gaussian densities with the same mean and variance (i.e., their Black-Scholes counterparts). All densities are plotted in terms of standardized (log) returns. By now it is well understood that options data with negatively sloped implied volatility surfaces, as in equations (5.2) and (5.3), correspond to negatively skewed and leptokurtic SPDs q_t (see Aït-Sahalia and Lo (1998)).

For comparison, we compute kernel estimates of the empirical marginal densities p_t using the historical returns data. Figure 2 presents the results for the same horizons as in Figure 1 above. Just as we constrained the option-implied SPDs for the log-returns to be Gaussian at long horizons, we constrain the empirical densities to also be Gaussian at the same horizons of 2.5 and one year for stocks and bonds, respectively. Since the optimal solution for a horizon t depends on the ratio q_t/p_t at that horizon, the impact of this constraint is limited to the constant of proportionality I_0 , which in turn does not distort the optimal choice because it affects all states and dates equally.

There are two striking differences between the two sets of densities. First, while both exhibit pronounced differences from normal distributions (see the second row of plots), the non-normalities of the empirical densities are much more concentrated toward short horizons than for the SPDs. This is especially the case for stocks, where the distribution of one-year returns is roughly as non-normal as the distribution of one-month returns. Empirically, in contrast, one-year returns are nearly Gaussian. Second, consistent with the literature, both

¹⁰Alternatively, a fully nonparametric fit is possible even with a single day's worth of option data if the appropriate model-free no-arbitrage shape restrictions are imposed, see Aït-Sahalia and Duarte (2003).

SPDs are considerably more negatively skewed than the empirical densities.

The next step is to use the Plackett formula to combine the marginal densities into joint densities. This requires that we first calibrate the coefficient θ in equation (3.13) to the estimated correlation of the stock and bond returns at different horizons t for each set of marginal densities. Table 3 gives details on this calibration. The second column of the table reports the empirical correlation between stock and bond returns at horizons ranging from one day to one year. The remaining columns provide the values of the coefficient θ for which the correlation implied by the resulting Plackett density matches the empirical correlation. Columns three, four, and five correspond to the option-implied SPDs q , the empirical densities p , and Gaussian densities p , respectively.

Figure 3 illustrates the relationship between the coefficient θ of the Plackett copula and the implied correlation ρ of stock and bond returns. The four lines in each plot correspond to return horizons of one month (solid), one quarter (dashed), one year (dashed-dotted), 2.5 years (dotted). The left plot is for the option-implied densities and the right plot is for the historical empirical densities. Since the marginal densities in both cases are constrained to be Gaussian at the 2.5-year horizon, the relationship between ρ and θ for Gaussian returns (at any horizon) is illustrated by the dotted line in either plot. The main message of this figure is that the calibrated values of θ are remarkably stable across different shapes of the marginal densities. As a result, our empirical results below are relatively insensitive to this intermediate calibration of the Plackett formula.

We can now combine these inputs to produce, at last, the objects of economic interest. Figures 4 and 5 present contour plots at horizons of one month, one quarter, one year, and 2.5 years of the joint option-implied and historical empirical densities, respectively. The joint densities are constructed by combining the marginal densities through the Plackett formula. Each contour corresponds to 10 percent cumulative probability. As expected from the marginal densities in Figure 1 and 2, the joint option-implied densities in Figure 4 exhibit substantially more non-normality than the empirical ones in Figure 5.

5.3 Relative Prices of Consumption

Given estimates of the option-implied SPDs q_t at different horizons, shown in Figure 4, we can use equation (2.25) to compute the consumption and portfolio rules of a CRRA investor for the various choices of the physical state density p_t , one of them being the historical empirical density in Figure 5. As we discussed in Section 2.4, the solution can be interpreted in two parts. First, the investor wants to consume more in states in which the price of consumption is cheap relative to the probability of realizing these states. Second, the investor wants to smooth consumption across states. In this section, we first examine the relative price effect, which is independent of the investor's preferences and is captured by the ratio q_t/p_t . In Section 5.4, we then examine the smoothing effect, which for CRRA preferences depends also on the coefficient of relative risk aversion γ .

Before exploring other possibilities, we present in Figure 6 as benchmark the ratio q_t/p_t under log-normality, corresponding to the Black-Scholes economy discussed above. The four plots show the relative prices of consumption in one month, one quarter, one year, and 2.5 years into the future. The SPD q_t is log-normal with moments matching the option-implied SPD. The physical distribution p_t is equal to the log-normal SPD shifted by risk premia that imply annualized Sharpe ratios of 0.5 for stocks and 0.05 for bonds (roughly in line with the historical moments reported in Table 2). We use these particular risk premia beliefs here and below not to advocate them as the best forecasts of future excess returns, but rather in order to focus the comparison between different cases on the shapes of the densities, rather than on their location.¹¹ In this figure and the following, we classify states (i.e., the joint realizations of returns for stocks and bonds) in terms of number of standard deviations from their respective SPD means.

The shading in the plot signifies the relative prices of consumption, with the black area as the most expensive states and the white area as the cheapest states. The legends next to each plot provide a rough scale of the relative price differences. For example, at the one month horizon, consumption in the black states is approximately $1.16/0.86 = 1.34$ times as

¹¹In fact, most estimates of the current equity risk premium are well below the historical equity risk premium. Fama and French (2001) estimate the equity risk premium to lie between 2.5% and 4.3%. Ibbotson and Chen (2003) estimate it to be between 4.0% and 6.0% at long horizons.

expensive as consumption in the white states.

To get a more precise reading of the relative prices of consumption, we tabulate in Table 4 the ratio q_t/p_t for nine states (all permutations of the bond and stock returns equal to their means and to their means plus or minus 1.5 standard deviations). For each state, the table shows a block of four numbers. The results for the benchmark case of log-normality are shown in the first row of each block.

At least three broad patterns emerge from Figure 6 and Table 4. First, consumption is most expensive in states associated with negative bond and stock returns, and it is cheapest in states associated with positive bond and stock returns. This pattern is consistent with the fact that both assets demand a positive risk premium. Second, the relative price gradient is considerably steeper across stock return states than bond return states, in line with the equity risk premium being substantially higher than the bond risk premium. Third, the spread between the cheapest and most expensive consumption states increases with the return horizon. This pattern is generated by the fact that both risk premiums increase approximately linearly with the return horizon (the effect is not exactly linear because of the different annualized risk premiums across horizons shown in Table 2).

Instead of assuming log-normality, we now combine the option-implied SPD q_t with each of the three different physical state densities p_t discussed in Section 4. We first construct p_t by shifting the option-implied SPD by risk premia that imply annualized Sharpe ratios of 0.5 for stocks and 0.05 for bonds. Figure 7 and the second row in each block of numbers in Table 4 report the ratios of the two densities, q_t/p_t , for this case. Notice that the results for the 2.5 year horizon are the same as for log-normality since we force the option-implied state price density to be log-normal beyond horizons for which options data is available.

The three broad patterns we discussed for the log-normal benchmark case are also clearly apparent with the option-implied distributions. There are, however, at least two important differences in the results. First, with the option-implied distributions the ratio q_t/p_t is lower for extreme stock return states and higher for extreme bond return states. The difference in q_t/p_t for the option implied distribution relative to log-normality is most pronounced for extreme positive stock return states, which means that the differences between extreme states increases considerably.

The second important difference is that with the option-implied distributions the relative prices of consumption for extreme positive stock return states (e.g., states beyond 1.5 standard deviations above the mean) become virtually insensitive to the realization of the bond return state, especially at intermediate horizons of one quarter to one year. This effect is best seen through the virtually horizontal contours toward the top of the quarterly and annual plots in Figure 7, as compared to the same plots in Figure 6.

Figures 8 and 9 present the density ratios q_t/p_t when p_t is either the empirical distribution of log-returns, scaled to match the option-implied volatility of q_t , or a log-normal distribution with the same moments, respectively. To illustrate the role of the risk premia, we consider in the first row risk premia that set the Sharpe ratios of stocks to 0.25 and that of bonds to 0.025. The second row corresponds to our benchmark Sharpe ratios of 0.5 and 0.05. Finally, in the third row of the figures the risk premia are such that the Sharpe ratios of stocks is 0.75 and that of bonds is 0.075. The third and fourth rows of Table 4 report the ratios q_t/p_t for our benchmark Sharpe ratios (0.5 for stocks and 0.05 for bonds).

Two important insights can be drawn from figures 8 and 9. First, the relative prices of state-contingent consumption are far less regular when the shape of the physical density is different from that of the state-price density. Since the physical densities all have the same first and second moments, the differences in results across figure 7-9 are entirely driven by higher-order moments. This means that the skewness and kurtosis of the investor's subjective beliefs about the likelihood of future states have first-order implications for intertemporal consumption and portfolio choice. Given the difficulties a typical investor has in forming beliefs about even the first two moments of returns, this result demonstrates the benefits of inferring the shape of the physical densities directly from options market data.

The second insight is that the prices of state-contingent consumption are relatively insensitive to variation in the subjective Sharpe ratios of stocks and bonds, especially at short horizons of one month or one quarter. This can be seen from the similarities of the shaded areas across rows of each figure. The shaded areas in the top row, corresponding to Sharpe ratios of 0.25 for stocks and 0.025 for bonds, are remarkably similar to those in the bottom row, corresponding instead to Sharpe ratios of 0.75 for stocks and 0.075 for bonds. This result is particularly striking in light of the sensitivity of the state prices to the higher moments of

the physical densities (comparing Figure 7 to the middle rows of figures 7-9).

5.4 Optimal Consumption Policies for Different p

Given the ratio q_t/p_t determined above, we now compute the optimal consumption plan, according to equation (2.25), across different realizations of the joint returns on stocks and bonds. Once the optimal consumption plan is determined, the investor's optimal portfolio strategy consists of purchasing a continuum of pure Arrow-Debreu securities. At date 0, the investor purchases for each state P and each date t a quantity $C_t^*(P) dt$ of the Arrow-Debreu security paying \$1 if $P_t = P$ and 0 otherwise. While not directly traded on exchanges, such Arrow-Debreu securities can be synthesized exactly using traded European call options on the underlying assets, or simply approximated in the form of butterfly payoffs (which converge to the Arrow-Debreu payoff if the strikes used to form the butterfly converge), as described in Section 2.6. In other words, the optimal investment strategy is fully characterized once we have determined the optimal consumption path.

As a benchmark for comparison, Figure 10 plots the optimal consumption of a CRRA investor under log-normal state-price and physical densities (the Black-Scholes case) with a risk aversion coefficient γ ranging from 2 to 10. The plot is constructed by applying equation (2.25) to the ratio q_t/p_t plotted in Figure 6. Figure 11 reports the same results when q_t is option-implied and p_t is mean-shifted from q_t by risk premia that imply a Sharpe ratio of 0.5 for stocks and 0.05 for bonds, corresponding to the ratio q_t/p_t plotted in Figure 7.

Figures 10 and 11 illustrate the trade-off between exploiting differences, across states, in the relative cost of state-contingent consumption and the desire to smooth consumption across states induced by risk aversion. As our formulation of the intertemporal problem makes clear, this trade-off is the economic essence of optimal consumption and portfolio choice. Other things equal, the investor likes to consume more in states in which the price of consumption is cheap relative to the (subjective) probability of realizing those states. That is, the optimal solution selects a higher C_t^* when the ratio q_t/p_t is low. Counteracting this effect, however, deviations across states from a constant value $C_t^* = C^*$ are penalized by the investor's risk aversion. The higher γ , the less sensitive C_t^* becomes to variations in q_t/p_t . In the limit,

as $\gamma \rightarrow \infty$, the optimal policy becomes $C_t^* = C^*$ irrespectively of the relative prices of consumption across states.

The smoothing effect is quite apparent as we move in either figure from the left column with $\gamma = 2$ to the right one with $\gamma = 10$. As the investor's risk aversion increases, the regions of (approximate) constancy of consumption across states, represented by the same shade of gray, increase. At the same time, the difference in shades of gray, from lightest to darkest, decreases. Combined, these patterns imply that as risk aversion increases, more states are associated with the same level of consumption and that the differences between consumption in the most extreme states diminishes. Risk aversion induces the investor to smooth consumption across states.

6. Conclusions

We developed an empirical approach to solving the intertemporal consumption and portfolio choice problem using option-implied information to determine the cost of consuming in, or equivalently betting against, future states of the world. Using the martingale representation method, we reduced the dynamic problem to a static one, in which the investor simply determines the level of state-contingent consumption in each state and at each date in the future. Our method gives a direct role to the investor's beliefs about the likelihood of future states, and we proposed different ways of implementing this subjective aspect of the problem.

Our results illustrate explicitly the inherent tension between, on the one hand, the investor's desire to exploit cross-sectional differences in state prices by consuming more in states in which the price of consumption is cheap relative to the probability of realizing these states and, on the other hand, the desire to smooth consumption across states induced by risk aversion. We argued that, as our formulation of the intertemporal problem makes clear, this trade-off is the economic essence of optimal consumption and portfolio choice.

Our method is potentially quite general. Unlike the dynamic programming approach to the intertemporal consumption and portfolio choice problem, our solution is always obtained explicitly so the numerical aspects of finding the solution are straightforward. One limitation of our method, however, is the need to be able to map future states of the world into

today's cross-section of option-implied state prices. This requires that we have access to the corresponding options data, which is a definite limitation especially for less standard asset classes. That said, while we have implemented the method for the typical cash/stocks/bonds asset allocation problem, with states of the world defined by the stochastic variation in the stock and bond indices, it is certainly possible to consider more complex asset allocation problems in which the investment opportunity set further varies due to randomness in the interest rate and/or stochastic volatility. Readily available data on interest rate options and variance swaps, respectively, make it possible to do so with our method.

References

- Aït-Sahalia, Y., and M. Brandt, 2001, “Variable Selection for Portfolio Choice,” *Journal of Finance*, 56, 1297–1351.
- Aït-Sahalia, Y., and A. Lo, 1998, “Nonparametric Estimation of State-Price-Densities Implicit in Financial Asset Prices,” *Journal of Finance*, 53, 499–547.
- Aït-Sahalia, Y., Y. Wang, and F. Yared, 2001, “Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?,” *Journal of Econometrics*, 102, 67–110.
- Banz, R. W., and M. H. Miller, 1978, “Prices for State-Contingent Claims: Some Estimates and Applications,” *Journal of Business*, 51, 653–672.
- Breeden, D. T., and R. H. Litzenberger, 1978, “Prices of State-Contingent Claims Implicit in Option Prices,” *Journal of Business*, 51, 621–651.
- Cox, J. C., and C.-F. Huang, 1989, “Optimum Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process,” *Journal of Economic Theory*, 49, 33–83.
- , 1991, “A Variational Problem Occurring in Financial Economics,” *Journal of Mathematical Economics*, 20, 465–487.
- Cvitanić, J., and I. Karatzas, 1995, “Hedging and Portfolio Optimization under Transaction Costs: Martingale Approach,” *Mathematical Finance*, 6, 133–165.
- Detemple, J., R. Garcia, and M. Rindisbacher, 2003, “A Monte Carlo Method for Optimal Portfolios,” *Journal of Finance*, 58, 401–446.
- Fama, E. F., and K. R. French, 1988, “The Early Exercise of Options on Treasury Bond Futures,” *The Journal of Financial and Quantitative Analysis*, 23, 437–449.
- , 2001, “The Equity Premium,” working paper, University of Chicago.
- Graham, J. R., and C. R. Harvey, 2002, “Expectations of Equity Risk Premia, Volatility and Asymmetry from a Corporate Finance Perspective,” working paper, Duke University.
- Harrison, M. J., and D. M. Kreps, 1979, “Martingales and Arbitrage in Multiperiod Securities Markets,” *Journal of Economic Theory*, 2, 381–408.
- Ibbotson, R. G., and P. Chen, 2003, “Long-Run Stock Returns: Participating in the Real Economy,” *Financial Analyst Journal*, Jan/Feb, 88–98.
- Karatzas, I., J. P. Lehoczky, and S. E. Shreve, 1987, “Optimal Portfolio and Consumption Decisions for a Small Investor on a Finite Horizon,” *SIAM Journal of Control and Optimization*, 25, 1557–1586.
- Merton, R. C., 1971, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model,” *Journal of Economic Theory*, 3, 373–413.

- , 1992, *Continuous Time Finance*. Basil Blackwell, New York, N.Y.
- Plackett, R. L., 1965, “A Class of Bivariate Distributions,” *Journal of the American Statistical Association*, 60, 516–522.
- Pliska, S. R., 1986, “A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios,” *Mathematics of Operations Research*, 11, 239–246.
- Rosenberg, J. V., 2003, “Nonparametric Pricing of Multivariate Contingent Claims,” *Journal of Derivatives*, 10, 9–26.
- Wachter, J. A., 2002, “Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets,” *Journal of Financial and Quantitative Analysis*, 37, 63–91.
- Wand, M. P., and M. C. Jones, 1995, *Kernel Smoothing*. Chapman and Hall, London, U.K.
- Welch, I., 2000, “Views of Financial Economists on the Equity Premium and on Professional Controversies,” *Journal of Business*, 73, 501–538.

Table 1: Options data

	Mean	Std Dev	Min	10%	50%	90%	Max
Stocks:							
Years to Maturity	0.8934	0.6108	0.0959	0.1753	0.7014	1.9260	1.9479
Moneyness	0.9506	0.1719	0.5367	0.7535	0.9481	1.1647	1.4970
Implied Volatility	0.2677	0.0426	0.1683	0.2167	0.2597	0.3284	0.4110
Bonds:							
Years to Maturity	0.3156	0.1837	0.0959	0.1178	0.2548	0.6767	0.7014
Moneyness	1.0507	0.0621	0.9080	0.9719	1.0446	1.1375	1.2249
Implied Volatility	0.1056	0.0318	0.0289	0.0616	0.1074	0.1389	0.2479

The table shows sample statistics for the options data on the S&P 500 index (stocks) and 10-year Treasury futures (bonds). The sample period is January 6, 2003 through January 17, 2003. The options prices are bid-ask midpoints recorded each day at 11:30am. There are 1093 observations for the stock options and 600 observations for the bond options.

Table 2: Underlying asset data

	Daily	Weekly	Monthly	Quarterly	Annual
Stocks:					
Mean	0.0793	0.0791	0.0794	0.0813	0.0829
Volatility	0.1462	0.1586	0.1494	0.1466	0.1447
Skewness	-1.6958	-1.6090	-0.7071	-0.6098	-0.5969
Kurtosis	47.7739	27.6679	7.1456	5.3098	3.5901
Bonds:					
Mean	0.0055	0.0050	0.0051	0.0062	0.0097
Volatility	0.1266	0.1385	0.1487	0.1531	0.1562
Skewness	-0.1132	-0.2310	-0.3242	-0.3297	-0.3094
Kurtosis	8.0242	6.7049	4.6273	3.6884	2.7549
Correlation	-0.2262	-0.2469	-0.2526	-0.2892	-0.4004

The table shows sample statistics for the excess returns on stocks and bonds over horizons ranging from one day to one year. Stocks and bonds represent the S&P 500 index and the on-the-run 10-year Treasury note, respectively. Returns are measured in excess of a maturity-matched risk-free zero-coupon yield. The sample period is January 2, 1962 through January 17, 2003.

Table 3: Calibrating the Plackett formula to historical correlations

Horizon	Calibrated θ			
	Historical Correlation	Option-Implied SPD q_t	Physical Density p_t	
			Empirical	Gaussian
Daily	-0.226	0.580	0.583	0.561
Weekly	-0.247	0.550	0.551	0.530
Monthly	-0.253	0.531	0.540	0.522
Quarterly	-0.289	0.446	0.473	0.464
Annual	-0.400	0.245	0.275	0.280

This table shows the historical correlations between stock and bond returns at different horizons. It also shows the values of the parameter θ for which the correlation implied by the Plackett formula for the option-implied SPDs and for the empirical or Gaussian physical distributions match the corresponding historical correlations.

Table 4: Relative prices of state-contingent consumption

Bond Return	<u>Stock Return</u>			<u>Stock Return</u>		
	-1.5 StdDev	Mean	+1.5 StdDev	-1.5 StdDev	Mean	+1.5 StdDev
	1 Month Horizon			1 Quarter Horizon		
-1.5 StdDev	2.480	0.818	0.254	1.984	1.119	0.559
	3.233	1.699	0.208	1.898	1.595	0.284
	0.817	1.003	0.422	0.791	0.975	0.367
	1.202	1.035	0.650	1.094	0.923	0.478
Mean	1.343	0.414	0.166	1.581	0.745	0.428
	2.368	0.869	0.135	1.711	1.105	0.223
	0.960	1.162	0.488	0.936	1.159	0.399
	0.933	0.895	0.568	1.079	0.930	0.444
+1.5 StdDev	0.696	0.283	0.123	1.056	0.568	0.370
	1.036	0.405	0.084	1.013	0.637	0.169
	0.800	0.953	0.402	0.806	0.925	0.314
	1.088	1.077	0.649	1.133	0.987	0.465
	1 Year Horizon			2.5 Year Horizon		
-1.5 StdDev	1.488	1.261	0.810	1.457	1.254	0.900
	1.361	1.417	0.319	-	-	-
	1.185	1.143	0.368	-	-	-
	1.171	1.096	0.350	-	-	-
Mean	1.445	0.969	0.678	1.433	1.098	0.744
	1.326	1.097	0.277	-	-	-
	1.251	1.047	0.214	-	-	-
	1.322	1.055	0.212	-	-	-
+1.5 StdDev	1.178	0.556	0.659	1.283	0.854	0.728
	1.073	0.846	0.271	-	-	-
	0.954	0.830	0.206	-	-	-
	1.010	0.797	0.201	-	-	-

This table shows the relative prices of state-contingent consumption, measured as the ratio of the SPD and physical distribution q_t/p_t , for a one month, one quarter, one year, and 2.5 year horizon. There are nine states comprised of stock and bond returns equal to their mean or their mean plus or minus 1.5 standard deviations (StdDev). For each state, the table shows four rows of numbers. In the first row, the SPD is log-normal with moments matching the option-implied SPD and the physical distribution is the log-normal SPD shifted by the historical risk premia. In the second row, the SPD is inferred from option prices and the physical distribution is the option-implied SPD shifted by the historical risk premia. In the third row, the SPD is inferred from option prices and the physical distribution is log-normal with moments matching the option-implied SPD but shifted by the historical risk premia. In the fourth row, the SPD is inferred from option prices and the physical distribution is the empirical distribution scaled to have the same moments as the option-implied SPD but shifted by the historical risk premia.

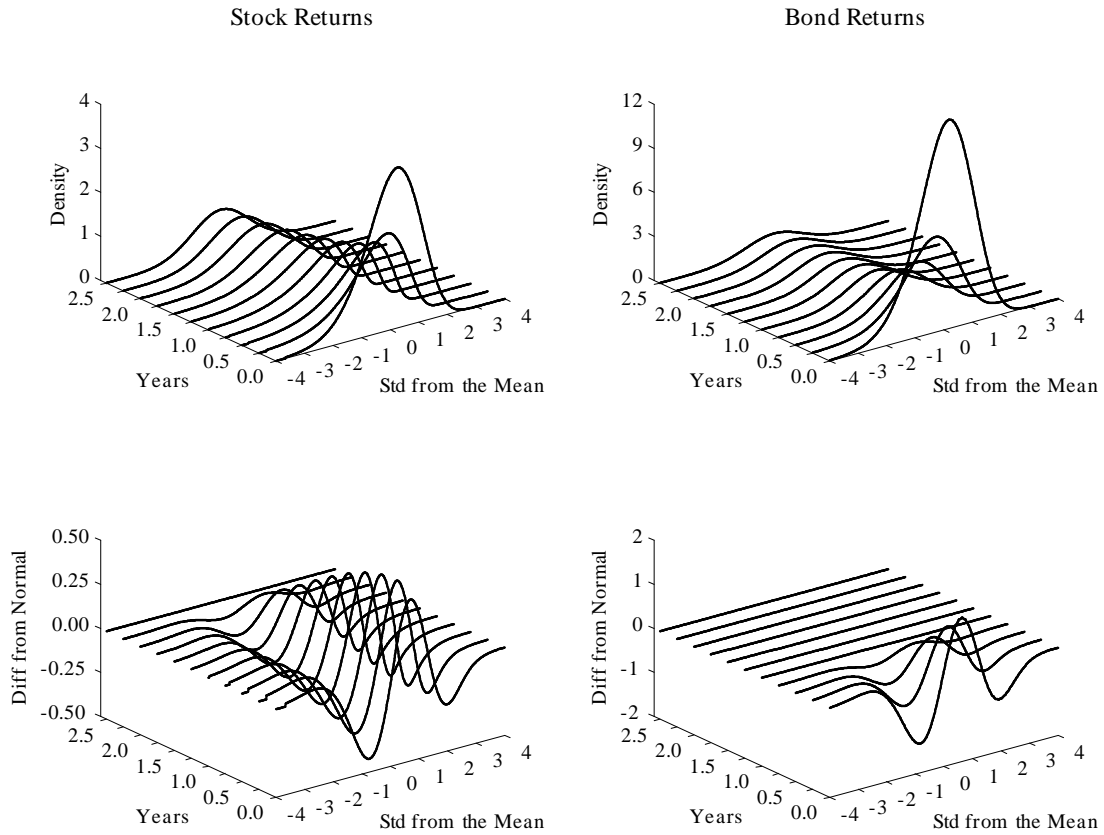


Figure 1
Marginal option-implied SPDs.

The first row of each plot shows the marginal option-implied SPDs of stock and bond returns. The second row shows the difference between these distributions and normal distributions with the same mean and variance.




Figure 2
Marginal empirical distributions.

The first row of each plot shows the marginal empirical distributions of stock and bond returns. The second row shows the difference between these distributions and normal distributions with the same mean and variance.




Figure 3
Calibrating the Plackett formula to correlations.

The plots show the relationship between the parameter θ of the Plackett formula and the implied correlation of stock and bond returns. The left plot is for the option-implied SPDs and the right plot is for the empirical distributions. The solid, dashed, dashed-dotted, and dotted lines correspond to a one month, one quarter, one year, and 2.5 year horizon, respectively. In the case of Gaussian densities, the relationship between ρ and θ is independent of the horizon and is equal to the 2.5 year case for either the option-implied or empirical distribution (since the 2.5-year distributions are constrained to be Gaussian in all cases).

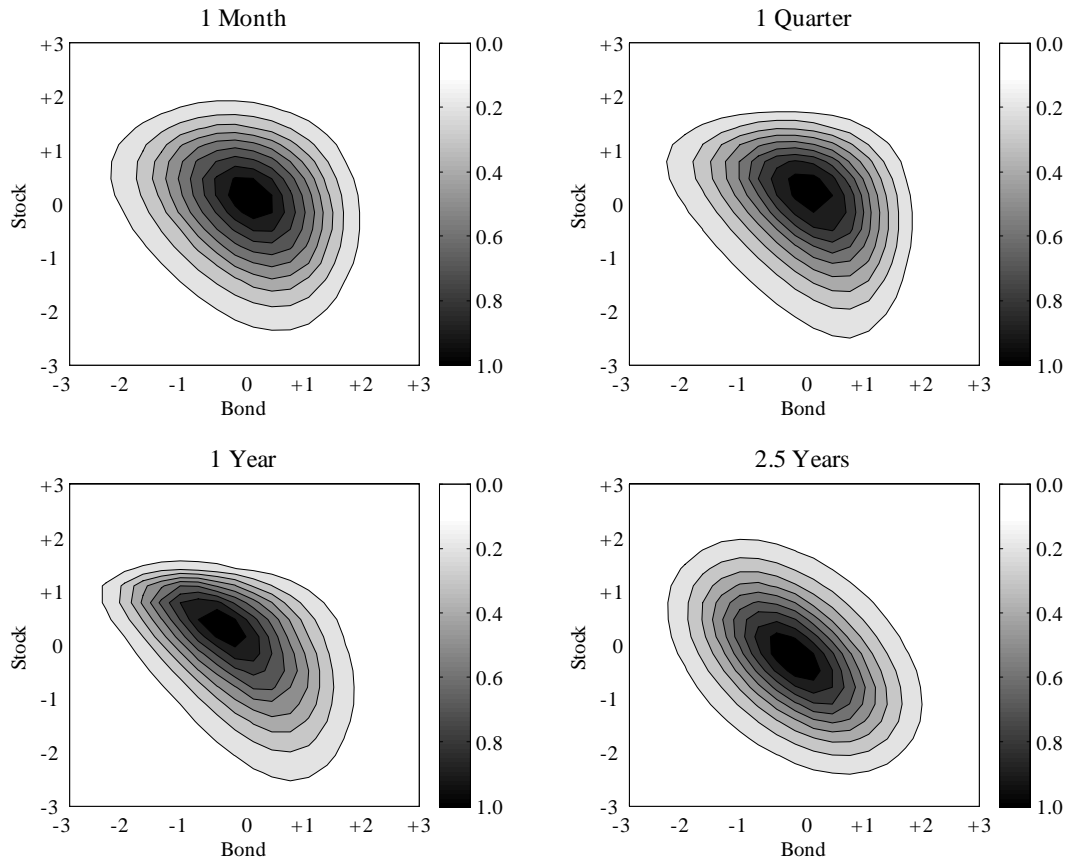


Figure 4
Joint option-implied SPDs.

The plots show the bivariate option-implied SPDs of stock and bond returns for a one month, one quarter, one year, and 2.5 year horizon. Each contour corresponds to 10% cumulative probability.




Figure 5
Joint empirical distribution.

The plots show the bivariate option-implied SPDs of stock and bond returns for a one month, one quarter, one year, and 2.5 year horizon. Each contour corresponds to 10% cumulative probability.




Figure 6
Relative prices of state-contingent consumption under log-normality.

The plots show the relative prices of state-contingent consumption, measured as the ratio of the SPD and physical distribution q_T/p_T , for a one month, one quarter, one year, and 2.5 year horizon. The SPD is log-normal with moments matching the option-implied SPD. The physical distribution is the log-normal SPD shifted by risk premia that imply an annualized Sharpe ratio of 0.5 for stocks and 0.05 for bonds.




Figure 7
Relative prices of state-contingent consumption with option-implied distributions.

The plots show the relative prices of state-contingent consumption, measured as the ratio of the SPD and physical distribution q_T/p_T , for a one month, one quarter, one year, and 2.5 year horizon. The SPD is inferred from option prices. The physical distribution is the option-implied SPD shifted by risk premia that imply an annualized Sharpe ratio of 0.5 for stocks and 0.05 for bonds.




Figure 8
Relative prices of state-contingent consumption with option-implied SPD and empirical physical distribution.

The plots show the relative prices of state-contingent consumption, measured as the ratio of the SPD and physical distribution q_T/p_T , for a one month, one quarter, and one year horizon. The SPD is inferred from option prices. The physical distribution is the empirical distribution scaled to have the same second moment as the option-implied SPD. The three rows of plots correspond to different beliefs about the Sharpe ratio of stocks and bonds. The means of the physical distribution are set to imply Sharpe ratios of 0.25 and 0.025 in the first row, 0.5 and 0.05 in the second row, and 0.75 and 0.075 in the third row.




Figure 9
Relative prices of state-contingent consumption with option-implied SPD and log-normal physical distribution.

The plots show the relative prices of state-contingent consumption, measured as the ratio of the SPD and physical distribution q_T/p_T , for a one month, one quarter, one year, and 2.5 year horizon. The SPD is inferred from option prices. The physical distribution is a log-normal density scaled to have the same second moment as the option-implied SPD. The three rows of plots correspond to different beliefs about the Sharpe ratio of stocks and bonds. The means of the physical distribution are set to imply Sharpe ratios of 0.25 and 0.025 in the first row, 0.5 and 0.05 in the second row, and 0.75 and 0.075 in the third row.




Figure 10
Smoothing state-contingent consumption across states under log-normality.

The plots show the exponentially scaled ratio of the SPD and physical distribution $(q_T/p_T)^{-1/\gamma}$ for a one month (first row), one quarter (second row), and one year horizons (third row). This scaled ratio is proportional to the optimal state-contingent consumption for an investor with CRRA utility and relative risk aversion γ . The SPD is log-normal with moments matching the option-implied SPD. The physical distribution is the log-normal SPD shifted by risk premia that imply an annualized Sharpe ratio of 0.5 for stocks and 0.05 for bonds.




Figure 11
Smoothing state-contingent consumption across states with option-implied distributions.

The plots show the exponentially scaled ratio of the SPD and physical distribution $(q_T/p_T)^{-1/\gamma}$ for a one month (first row), one quarter (second row), and one year horizons (third row). This scaled ratio is proportional to the optimal state-contingent consumption for an investor with CRRA utility and relative risk aversion γ . The SPD is inferred from option prices. The physical distribution is the option-implied SPD shifted by risk premia that imply an annualized Sharpe ratio of 0.5 for stocks and 0.05 for bonds.