Self-Control Preferences and the Volatility of Stock Prices

David DeJong*
University of Pittsburgh

Marla Ripoll
University of Pittsburgh

April, 2003

Abstract

We use a simple asset-pricing environment to study the ability of self-control preferences to account for the behavior of aggregate annual stock prices and dividends. Using as benchmarks results obtained using CRRA and habit/durability preferences, we find empirical support for a positive role for temptation in characterizing the behavior of these data. Moreover, parameter estimates obtained using a full-information procedure indicate a quantitatively significant presence of temptation. This presence generates an annual return premium of 1.32 percentage points relative to CRRA preferences, and a utility loss equivalent to a reduction in steady state consumption of 5.25%. While self-control preferences do not fully account for the level of price volatility observed in the data, our results are nevertheless encouraging, particularly in light of the simple environment upon which our analysis is based.

Keywords: full-information estimation; posterior analysis; risk aversion; habit/durability
JEL Codes: E44; C11; C15

*Corresponding author: DeJong; Department of Economics, University of Pittsburgh, Pittsburgh, PA, 15260; dejong@pitt.edu. For helpful comments we thank Per Krusell, Anthony A. Smith, Jr. and seminar participants at the Institute for Advanced Studies (Vienna), and Charles, Comenius and Sabanci Universities. Usual disclaimer applies.

1 Introduction

Following LeRoy and Porter (1981) and Shiller (1981), considerable effort has gone towards understanding why aggregate stock prices are so volatile relative to corresponding discounted dividend streams. An important component of this effort has focused on the role played by consumers' preferences in accounting for this behavior. For example, Grossman and Shiller (1981) explored whether movements in intertemporal marginal rates of substitution in consumption induced by constant-relative-risk aversion (CRRA) preferences are capable of generating observed patterns of stock-price volatility, but found these movements to be inadequate. And Heaton (1995) showed that while certain parameterizations of habit persistence preferences are capable of generating substantial increases in marginal rates of substitution, and thus have the potential to account for the puzzle, such parameterizations appear empirically implausible when seeking to account for asset-pricing behavior more generally (i.e., in expanding the set of moments beyond volatility measures to obtain parameter estimates).

Here, we examine empirically whether self-control preferences can plausibly account for observed interactions between prices and dividends, including the relative volatility of prices. The potential role of self control in influencing asset-pricing behavior is becoming increasingly recognized. For example, as noted by Gul and Pesendorfer (2000), the propensity for experimental subjects to willingly forgo future payoffs in return for relatively modest current rewards suggests this potential. (See Kocherlakota, 2001, for scepticism regarding this view.) If agents face temptations to deviate from intertemporally optimal savings plans in return for higher current consumption, this would carry implications for their demand for assets, and thus for asset prices. Our interest is in quantifying potential implications of these temptations.

Our analysis is based on three versions of a Lucas (1978) – type environment featuring a representative household and a single asset. The versions differ only in their specification of preferences. In each model, the price of the asset reflects the future dividend stream it is expected to generate, weighted by the household's marginal rate of substitution. Consumption is financed by dividend payments and an exogenous endowment; innovations to these stochastic processes drive stock prices.

In the data, stock prices are not only far more volatile than dividends, but their fluctuations are also closely correlated. Absent this correlation, virtually any preference specification could account for the observed pattern of volatility in this environment, if coupled with a sufficiently volatile endowment process. The problem is that endowment innovations weaken the link between price and dividend movements. Accounting for both the high volatility and correlation patterns observed in the data thus represents the crux of the empirical challenge facing the model.

The first specification we consider features CRRA preferences; the second features habit/durability preferences parameterized, e.g., as in Ferson and Constantinides (1991) and Heaton (1995); and the third features preferences specified following Gul and Pesendorfer (2000), under which the household faces a temptation to deviate from its intertemporally optimal consumption plan by selling its entire holding of shares and maximizing current-period consumption. The household never succumbs to this temptation in equilibrium, yet the presence of temptation potentially wields substantial influence over the household's demand for shares.

Our analysis proceeds in two steps. The first involves estimating the parameters of each of the environments using the full-information Bayesian methodology of DeJong, Ingram and Whiteman (2000). The second involves assessing the ability of each model to account for selected aspects of the time-series behavior of the data, as summarized by a set of moments. The data we use are the annual Standard & Poor's 500 stock price and dividend series analyzed by Shiller (1981), extended through 1999. The moments we evaluate are the standard deviations of prices and dividends; the ratio of these standard deviations; and the correlation between prices and dividends. Posterior distributions of the models' structural parameters are obtained by combining prior distributions with likelihood functions corresponding to linear approximations of the models. From these distributions, we construct corresponding posterior distributions of the set of moments used to evaluate empirical performance. Finally, we compare these distributions with counterparts obtained using a vector-autoregressive (VAR) model specified for dividends and prices; the VAR distributions serve as the summary of the data we use as a basis upon which to judge the models' empirical performance. Thus our estimates of the structural models reflect the full set of empirical implications conveyed by their associated likelihood functions and our *a priori* views regarding parameter values; and the subsequent analysis of fit reflects performance along a select subset of dimensions.

Our analysis of the models featuring CRRA and habit/durability preferences serves as a benchmark for interpreting the results we obtain under self-control preferences. Here, we obtain results consistent with previous studies: neither specification is capable of accounting for observed stock-price volatility while at the same time preserving the close correlation observed between prices and dividends. CRRA preferences turn out to provide relatively superior performance along the volatility dimension relative to habit/durability preferences, but do so at the cost of performing relatively poorly along the correlation dimension.

Relative to these specifications, we find that the self-control specification provides an improved characterization of the data, but not a full account of stock-price volatility. Regarding parameter estimates, the self-control specification yields CRRA preferences as a special case under a certain parameter restriction. Despite using a prior distribution with a modal value corresponding with this restriction, the posterior distribution is shifted distinctly away from this point, indicating empirical support for the presence of a temptation effect. The strength of the effect is significant: households endowed with the self-control preferences we estimate would demand a 1.32% annual return premium relative to households endowed with CRRA preferences. Moreover, the temptation cost borne by households is equivalent to a reduction in steady state consumption of 5.25%.

Regarding empirical performance, the self-control specification marginally outperforms the CRRA specification: a posterior-odds comparison rates the overall performance of the self-control specification as superior to the CRRA specification by a margin of approximately 2.4:1. Nevertheless, the self-control specification fails to account fully for the volatility of stock prices observed in the data, which is perhaps not surprising given the simple structure within which it is embedded. Thus we conclude in the context of this simple environment that the presence of temptation seems to provide an important step towards a better understanding of stock-price volatility, if not a complete resolution of the puzzle.

As noted, there is a growing body of work that focuses on the potential role played by temptation and other self-control problems in influencing asset-pricing behavior. Perhaps most closely related to our study is that of Krusell, Kuruscu and Smith (2002), who explore the ability of a model featuring a small subset of agents who face a temptation to save to account for the equity premium puzzle. These agents end up dominating the equity market, thus their behavior determines asset prices. Calibrating their model to match the wealth distribution in the U.S., Krusell et al. find that they can account for the equity premium and low risk-free rate with a risk-aversion parameter in the neighborhood of 22.

2 Preference specifications

As noted, we consider a Lucas (1978)- type environment featuring a representative household and a single asset. Time-$(t-1)$ share holdings s_{t-1} yield a dividend payment d_t at time t; time-t share prices are given by p_t. Households maximize expected lifetime utility by financing consumption c_t from dividend earnings, proceeds from sales of shares, and an exogenous endowment e_t:

$$\max_{c_t} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t),$$

$0 < \beta < 1$, subject to

$$c_t + p_t(s_t - s_{t-1}) = d_t s_{t-1} + e_t. \tag{1}$$

Since households are identical, in equilibrium $s_t = s_{t-1}$ for all t, and thus $c_t = d_t + e_t$ (hereafter s_t is normalized to 1). Combining this equilibrium condition with the household's necessary condition for a maximum yields the pricing equation

$$p_t = \beta E_t \frac{u'(d_{t+1} + e_{t+1})}{u'(d_t + e_t)} (d_{t+1} + p_{t+1}). \tag{2}$$

The model is closed by specifying stochastic processes for (d_t, e_t). These are given by

$$\ln d_t = (1 - \rho_d) \ln d^* + \rho_d \ln d_{t-1} + \varepsilon_{dt} \tag{3}$$

$$\ln e_t = (1 - \rho_e) \ln e^* + \rho_e \ln e_{t-1} + \varepsilon_{et}$$

with

$$\begin{bmatrix} \varepsilon_{dt} \\ \varepsilon_{et} \end{bmatrix} \sim iid \;\; N(0, \Sigma).$$

From (2), it is clear that following a shock to either d_t or e_t, the response of p_t depends in part upon the variation of the marginal rate of substitution between t and $t+1$. This in turn depends upon the momentary utility function $u(.)$. We now specify the alternative utility functions considered here, and evaluate their potential for generating relatively volatile responses of p_t.

2.1 CRRA

CRRA preferences are parameterized as

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \tag{4}$$

where $\gamma > 0$ measures the degree of relative risk aversion, and $1/\gamma$ is the intertemporal elasticity of substitution. In this case, the equilibrium pricing equation reads

$$p_t = \beta E_t \frac{(d_{t+1} + e_{t+1})^{-\gamma}}{(d_t + e_t)^{-\gamma}} (d_{t+1} + p_{t+1}). \tag{2a}$$

Notice that, *ceteris paribus*, a relatively large value of γ will increase the volatility of price responses to exogenous shocks, at the cost of decreasing the correlation between p_t and d_t.

2.2 Habit / durability

Following Ferson and Constantinides (1991) and Heaton (1995), the habit/durability specification we consider is parameterized as

$$u(x_t) = \frac{x_t^{1-\gamma}}{1-\gamma}, \tag{4b}$$

with

$$x_t = x_t^d - \alpha x_t^h,$$

where $0 < \alpha < 1$, x_t^d is the household's durability stock, and x_t^h its habit stock. The stocks are defined by

$$x_t^d = \sum_{j=0}^{\infty} \delta^j c_{t-j}$$

$$x_t^h = (1-\theta)\sum_{j=0}^{\infty} \theta^j x_{t-1-j}^d = (1-\theta)\sum_{j=0}^{\infty} \theta^j \sum_{i=0}^{\infty} \delta^i c_{t-1-i}$$

where $0 < \delta < 1$ and $0 < \theta < 1$. Thus, the durability stock represents the flow of services from past consumption, which depreciates at rate δ. This parameter also represents the degree of intertemporal consumption substitutability. The habit stock can be interpreted as a weighted average of the durability stock, where the weights sum to one. Notice that more recent durability stocks, or more recent flows of consumption, are weighted relatively heavily. Thus, the presence of habit captures intertemporal consumption complementarity. The variable x_t represents the current level of durable services net of the average of past services; the parameter α measures the fraction of the average of past services that is netted out. Notice that if $\delta = 0$, we would only have habit persistence, while if $\alpha = 0$ we would only have durability. Finally, when $\theta = 0$, the habit stock includes only one lag.

Using the definitions of durable and habit stocks, x_t becomes

$$x_t = c_t + \sum_{j=1}^{\infty} \left[\delta^j - \alpha(1-\theta) \sum_{i=0}^{j-1} \delta^i \theta^{j-i-1} \right] c_{t-j} \equiv \sum_{j=0}^{\infty} \Phi_j c_{t-j},$$

where we define $\Phi_0 \equiv 1$. Thus, for these preferences, the equilibrium pricing equation is given by

$$p_t = \beta E_t \frac{\sum_{j=0}^{\infty} \beta^j \Phi_j \left(\sum_{i=0}^{\infty} \Phi_i c_{t+1+j-i} \right)^{-\gamma}}{\sum_{j=0}^{\infty} \beta^j \Phi_j \left(\sum_{i=0}^{\infty} \Phi_i c_{t+j-i} \right)^{-\gamma}} (d_{t+1} + p_{t+1}) \tag{2b}$$

where as before $c_t = d_t + e_t$.

To see how the presence of habit and durability can potentially influence the volatility of the prices, rewrite the pricing equation as

$$p_t = \beta E_t \frac{(c_{t+1} + \Phi_1 c_t + \Phi_2 c_{t-1} + ...)^{-\gamma} + \beta\Phi_1(c_{t+2} + \Phi_1 c_{t+1} + \Phi_2 c_t + ...)^{-\gamma} + ...}{(c_t + \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + ...)^{-\gamma} + \beta\Phi_1(c_{t+1} + \Phi_1 c_t + \Phi_2 c_{t-1} + ...)^{-\gamma} + ...} (d_{t+1} + p_{t+1}) .$$

When there is a positive shock to say e_t, c_t increases by the amount of the shock, say σ_e. Given (3), c_{t+1} would increase by $\rho_e\sigma_e$, c_{t+2} would increase by $\rho_e^2\sigma_e$, etc. Let us examine the first term in parenthesis both in the numerator and the denominator. First, in the denominator c_t will grow by σ_e. Second, in the numerator $c_{t+1} + \Phi_1 c_t$ goes up by $(\rho_e + \Phi_1)\,\sigma_e \lesseqgtr \sigma_e$. Thus, whether the share price p_t increases by more than in the standard CRRA case depends ultimately on whether $\rho_e + \Phi_1 \lesseqgtr 1$. Notice that if $\Phi_j = 0$ for $j > 0$, then the equation above reduces to the standard CRRA utility case. If we had only habit and not durability, then $\Phi_1 < 0$, and thus the response of prices would be greater than in the CRRA case. This result is intuitive: habit captures intertemporal complementarity in consumption, which strengthens the smoothing motive relative to the time-separable CRRA case.

Alternatively, if we had only durability and not habit, then $0 < \Phi_1 < 1$, but we still would not know whether $\rho + \Phi_1 \lesseqgtr 1$. Thus, with only durability, we cannot conclude on how the volatility of p_t would be affected since it will depend upon the sizes of ρ and Φ_1. Finally, we also face indeterminacy under a combination of both durability and habit: if α is large and δ is small enough to make $\rho + \Phi_1 < 1$, then we would get increased price volatility. Thus this issue is fundamentally quantitative.

2.3 Self-control

We now consider a household that every period faces a temptation to consume all of its wealth. Resisting this temptation imposes a self-control utility cost. To model these preferences we follow Gul and Pesendorfer [2000], who identified a class of dynamic self-control preferences. In this case, the problem of the household can be formulated recursively as

$$W(s,P) = \max_{s'}\{u(c) + v(c) + \beta EW(s',P')\} - \max_{\widetilde{s}'} v(\widetilde{c})$$

where $P = (p, d, e)$; $u(.)$ and $v(.)$ are Von Neuman-Morgesten utility functions; $0 < \beta < 1$; $\widetilde{c}$ represents temptation consumption; and s' denotes share holdings next period. While $u(.)$ is the momentary utility function, $v(.)$ represents temptation. The problem above is subject to the following budget constraints

$$c = ds + e - p(s' - s)$$

$$\widetilde{c} = ds + e - p(\widetilde{s}' - s)$$

In the specification above, $v(c) - \max_{\widetilde{s}'} v(\widetilde{c}) \leq 0$ represents the disutility of self-control given that the agent has chosen c. The concavity/convexity of $v(.)$ turns out to carry important implications for stock-price volatility. We consider both scenarios in turn. (Notice that if $v(.)$ is convex, one must impose restrictions to guarantee that $u(.) + v(.)$ is strictly concave. On the other hand, if $v(.)$ is strictly concave, one must guarantee that $W(.)$ is also strictly concave. It turns out, as shown later, that this is the case.)

Throughout the paper, we will consider strictly increasing specifications for $v(.)$. In this case, the solution for $\max_{\widetilde{s}'} v(\widetilde{c})$ is simply to drive $\widetilde{c}$ to the maximum allowed by the constraint $\widetilde{c} = ds + e - p(\widetilde{s}' - s)$, which is attained by setting $\widetilde{s}' = 0$. Thus, we can rewrite the problem as

$$W(s, P) = \max_{s'} \{u(c) + v(c) + \beta E W(s', P')\} - v(ds + e + ps)$$

subject to

$$c = ds + e - p(s' - s).$$

The optimality condition reads

$$\left[u'(c) + v'(c)\right] p = \beta E W'(s', P')$$

and since

$$W'(s, P) = \left[u'(c) + v'(c)\right] (d + p) - v'(ds + e + ps)(d + p)$$

then the optimality condition becomes

$$\left[u'(c) + v'(c)\right] p = \beta E \left[u'(c') + v'(c') - v'(d's' + e' + p's')\right] (d' + p').$$

Combining this expression with the equilibrium conditions $s = s' = 1$ and $c = d + e$ yields

$$p = \beta E \left(d' + p'\right) \left[\frac{u'(d' + e') + v'(d' + e') - v'(d' + e' + p')}{u'(d + e) + v'(d + e)}\right].$$

Notice that when $v(.) = 0$, there is no temptation, and the equation above reduces to the standard case. Otherwise, the term $u'(d' + e') + v'(d' + e') - v'(d' + e' + p')$ represents tomorrow's utility benefit from saving today. This corresponds to the standard marginal utility of wealth tomorrow $u'(d' + e')$, plus the term $v'(d' + e') - v'(d' + e' + p')$ which represents the derivative of the utility cost of self-control with respect to wealth. The behavior of this term depends on the concavity/convexity of $v(.)$. For $v(.)$ convex, this derivative is negative, implying that the standard marginal utility of wealth tomorrow $u'(d'+e')$ is decreased by the term $v'(d'+e')-v'(d'+e'+p') < 0$. In other words, the convexity of $v(.)$ implies that as wealth increases, self-control becomes more costly.

We assume the following functional forms for the momentary and temptation utility functions

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \tag{4c}$$

$$v(c) = \lambda \frac{c^{\phi}}{\phi}$$

with $\lambda > 0$, which imply the following pricing equation:

$$p = \beta E \left[d' + p'\right] \left[\frac{(d' + e')^{-\gamma} + \lambda(d' + e')^{\phi-1} - \lambda(d' + e' + p')^{\phi-1}}{(d + e)^{-\gamma} + \lambda(d + e)^{\phi-1}}\right]. \tag{2c}$$

Continuing with the case in which $v(.)$ is convex, so that $\phi > 1$, a positive shock to d or e would imply *less* price volatility than under the CRRA case. To see this, rewrite (2c) as

$$p = \beta E\left[d' + p'\right]\left[\frac{\frac{(d'+e')^{-\gamma}}{(d+e)^{-\gamma}} + \lambda(d+e)^{\gamma}\left[(d'+e')^{\phi-1} - (d'+e'+p')^{\phi-1}\right]}{1+\lambda(d+e)^{\phi-1+\gamma}}\right]$$

where we can see that a larger, say e, increases the denominator, while decreasing the term $\lambda(d+e)^{\gamma}\left[(d'+e')^{\phi-1} - (d'+e'+p')^{\phi-1}\right]$ in the numerator. Both effects imply that relative to the CRRA case, in which $\lambda = 0$, this specification *reduces* price volatility in the face of an endowment shock, which is precisely the opposite of what we would like to achieve in order to better match the data.

The mechanism behind this reduction in price volatility is as follows: a positive shock to d or e increases the household's wealth today, which has three effects. The first, which we call "smoothing" captures the standard intertemporal motive: the household would like to increase saving, which drives up the share price. Second, there is a "temptation" effect: with more wealth today, the feasible budget set for the household increases, which represents more temptation to consume, and less willingness to save. This effect works opposite to the first, and reduces price volatility with respect to the standard case. Third, we have the "self-control" effect: due to the assumed convexity of $v(.)$, marginal self-control costs also increase, which reinforces the second effect. As shown above, the last two effects dominate the first, and thus under convexity of $v(.)$ the volatility is reduced relative to the CRRA case.

In contrast, price volatility would not necessarily be reduced if $v(.)$ is concave, and thus $0 < \phi < 1$. In this case, when d or e increases, the term $\lambda(d+e)^{\gamma}\left[(d'+e')^{\phi-1} - (d'+e'+p')^{\phi-1}\right]$ increases. On the other hand, if $\phi - 1 + \gamma > 0$, i.e., if the risk-aversion parameter $\gamma > 1$, the denominator also increases. If the increase in the numerator dominates that in the denominator, then we will observe higher price volatility than in the CRRA case.

To understand this effect, note that the derivative of the utility cost of self-control with respect to wealth is positive if $v(.)$ is concave: $v'(d'+e') - v'(d'+e'+p') > 0$. This means that as agents get wealthier, self-control costs become lower. This explains why it might be possible to get higher price volatility in this case. In fact, the mechanism behind this result still involves the three effects discussed above: smoothing, temptation, and self-control. The difference is on the latter effect:

under concavity, self-control costs are *decreasing* in wealth. This gives the agent an incentive to save *more* rather than less. If this self-control effect dominates the temptation effect, then these preferences will produce higher price volatility.

Notice that when $v(.)$ is concave, we need to impose conditions to guarantee that $W(.)$ is strictly concave, so that the solution corresponds to a maximum. In particular, the second derivative of $W(.)$ must be negative:

$$-\gamma\,(d+e)^{-\gamma-1}+\lambda(\phi-1)\left[(d+e)^{\phi-2}-(d+e+p)^{\phi-2}\right]<0$$

which in fact holds for any d, e, and $p>0$, and for $\gamma>0$, $\lambda>0$, and $0<\phi<1$. Hereafter, we will proceed under this set of parameter restrictions.

2.4 Steady States

We conclude with comments regarding steady states. As noted above, s is normalized to one. We also normalize the steady-state value $d^*=1$. Let $\eta=\frac{e^*}{d^*}$, so that $\eta=e^*$. Thus, $c^*=1+\eta$. From (2a) and (2b), the steady-state price under CRRA and habit/durability preferences is given by

$$p^*=\frac{\beta}{1-\beta}d^*=\frac{\beta}{1-\beta}. \tag{5}$$

From the optimality conditions under self-control preferences, steady-state temptation consumption is $\widetilde{c}^*=1+\eta+p^*$. From (2c), the steady-state price in this case is given by

$$p^*=\beta\,(1+p^*)\left[\frac{(1+\eta)^{-\gamma}+\lambda\,(1+\eta)^{\phi-1}-\lambda\,(1+\eta+p^*)^{\phi-1}}{(1+\eta)^{-\gamma}+\lambda\,(1+\eta)^{\phi-1}}\right]. \tag{6}$$

As we shall see below, (5) and (6) carry significant influence on parameter estimates: (5) exclusively for β, (6) for $(\beta,\phi,\gamma,\eta,\lambda)$. The interpretation of (5) is straightforward. Letting $\beta=1/(1+r)$, where r denotes the household's discount rate, (5) implies $p^*/d^*=1/r$. Thus as the household's discount rate increases, its asset demand decreases, driving down the steady state price level. Empirically, the average price/dividend ratio observed in the data serves to pin down β, under CRRA and habit/durability preferences.

Regarding (6), the left-hand-side is a 45-degree line. The right-hand side is strictly concave in p^*, has a positive intercept, and a positive slope that is less than one at the intercept. Thus (6) yields a unique positive solution for p^* for any admissible parameterization of the model. (In practice, we solve for p^* in (6) numerically using GAUSS's quasi-Newton algorithm NLSYS.) An increase in λ causes the function of p^* on the right-hand-side of (6) to shift down and flatten, thus p^* is decreasing in λ. The intuition for this is again straightforward: an increase in λ represents an intensification of the household's temptation to liquidate its asset holdings. This drives down its demand for asset shares, and thus p^*. Note the parallel between this effect and that generated by an increase in r, or a decrease in β, which operates analogously in both (5) and (6). The impact on p^* of the remaining parameters is discussed below.

3 Empirical Implementation

As noted, our empirical analysis consists of two components: model estimation, and evaluation of fit. The approach we use to evaluate fit is described in Section 4; estimation is accomplished using the full-information Bayesian procedure developed by DeJong, Ingram and Whiteman (2000). A sketch of the technical details of this procedure is provided in the appendix. Briefly, this procedure works as follows. First, we obtain a log-linear approximation of the relevant asset-pricing equation (2) using a first-order Taylor-series expansion. Coupled with the specifications of the exogenous dividend and endowment processes given in (3), this yields the likelihood function for dividends and prices implied by the model. Next, since we generally have clear *a priori* guidance from theory regarding values of the models' structural parameters, we bring this guidance to bear formally by specifying prior distributions over these parameters. Finally, combining these priors with the likelihood function using Bayes' Rule, we obtain posterior distributions of the parameters, along with posterior distributions of any functions of these parameters we wish to evaluate.

Under this procedure, the parameters and functions of parameters we analyze are interpreted as random variables. Their corresponding posterior distributions reflect the relative plausibility of alternative values of these variables, conditional on the models, our priors, and the observed data. Focusing on parameters, their posterior distributions convey information, e.g., regarding the extent to which the data support the presence of a non-trivial temptation effect under the self-control specification. Focusing on a function such as the standard deviation of prices, its posterior distributions convey information regarding the extent to which the restricted stochastic specifications of prices implied by the alternative structural models we consider coincide with the actual behavior of prices, as characterized by the stochastic specification associated with a relatively unrestricted reduced-form representation such as a VAR.

In specifying prior distributions, our goal is to place an emphasis on standard ranges of parameter values, while assigning sufficient prior uncertainty to allow the data to have a nontrivial influence on our estimates. The priors we use for this purpose are summarized in Table 1. In all cases, prior correlations across parameters are zero; and each of the informative priors we specify is normally distributed (but truncated when appropriate).

Consider first the parameters common across all models. The prior mean (standard deviation) of the discount factor β is 0.96 (0.02), implying an annual discount rate of 4%. The prior for the risk aversion parameter γ is centered at 2 with a standard deviation of 1. We specified the usual non-informative prior over the covariance matrix of the shocks Σ, proportional to $det(\Sigma)^{-(m+1)/2}$ (where m is the dimension of Σ), and centered the priors over the persistence parameters (ρ_d, ρ_e) at 0.9, with standard deviations of 0.05. The prior over $\eta = e^*/d^*$ is more problematic: we wish to use this prior to de-emphasize extremely large and small values of this ratio, but have little guidance upon which to sharply judge what constitutes what such values might be. So we considered two alternative specifications of prior means (standard deviations) for η: 5 (3); 10 (5). It turns out that the data are relatively uninformative regarding the location of η, so that the prior dominates the posterior along this dimension; this is particularly true in the CRRA specification. Fortunately, the prior over η turns out have little influence either on the fit of the models or on estimates of the additional parameters, since posterior correlations between η and the additional parameters are limited to a single dimension ($\sigma_{\varepsilon e}$, for reasons discussed in Section 4). Given the limited general sensitivity of our results to the prior over η, we report in detail only those results obtained using the (10, 5) specification. Finally, the priors over (β, ρ_d, ρ_e) are truncated from below at zero and from above at one, and the prior over η is truncated from below at zero.

Regarding the habit/durability parameters, we centered the priors over the decay parameters δ and θ at 0.0625, and assigned standard deviations of 0.03. This follows the prior specification employed by Otrok (2001), who centered his priors at 0.5 in working with quarterly data. This specification for δ is also consistent with SMM estimates obtained by Heaton (1995), while the specification for θ lies considerably below his estimates of this parameter (although as shown below, the posterior mode of θ turns out to be zero in this setting, so our prior does not account for the difference in results we obtain along this dimension). Finally, we specified an uninformative prior over α, which indicates the relative importance of habit and durability in the preference specification; this was done because we have no *a priori* reason to restrict attention to a particular range of the parameter space in this case. In contrast, Otrok centered his prior over α at 0.5; as shown below, the adoption of a similar prior in this setting would have a trivial influence on our results. All three priors were truncated from below at zero and from above at 1.

Regarding the self-control parameters λ and ϕ, we specified priors by focusing on the implications these parameters carry for the steady state relationship between prices and dividends, as indicated in (6). Recall that along with λ and ϕ, the parameters β, γ and η also influence this relationship. When $\lambda = 0$, we revert to the CRRA case, under which our prior over β is centered on 0.96, implying an annual rate of return of 4%, or a steady state price/dividend ratio of 24. As we shall see below, the average rate of return in the data is 4.49%, implying an average price/dividend ratio of 22.26. Thus we deem parameterizations of λ and ϕ that leave us in this ballpark as reasonable *a priori*. Fixing β, γ and η at their prior means, (λ, ϕ) combinations in the respective ranges ([0, 0.001], [0, 0.75]) deliver this behavior. Increasing β by one prior standard deviation, up to 0.98, moves these ranges to ([0, 0.006], [0, 0.8]); re-centering β at 0.96 and decreasing γ by one standard deviation to 1 moves these ranges to ([0, 0.014], [0, 0.8]); and re-centering γ and decreasing η by one standard deviation to 5 moves these ranges to ([0, 0.004], [0, 0.7]). Thus notice that small changes in λ can have large impacts on steady state price/dividend values, while changes in ϕ have relatively small impacts. Recall the intuition behind the impact of λ: as the household's temptation to consume increases, its demand for asset holdings decreases, thus driving down p^*. In light of these calculations, we specified a normal distribution for λ centered and truncated from below at 0, with a standard deviation of 0.01; and a normal distribution for ϕ centered at 0.4 with a standard deviation of 0.2, truncated from below at 0 and from above at 1.

Since the model approximations we work with involve logged deviations of variables from steady state values, the data used to analyze the models must be transformed accordingly. Shiller's (1981) analysis of stock-price volatility was based on the assumption that dividends and prices are trend stationary; DeJong (1992) and DeJong and Whiteman (1991, 1994) provided subsequent empirical support for this assumption. Following this work, we too adopt the trend-stationarity assumption, which implies that logged steady state dividends follow a linear trend. This assumption, coupled with the steady state specifications for prices and dividends given either by (5) or (6), implies a restricted trend-stationarity specification for prices. Specifically, given the value of p^*/d^* implied under either (5) or (6), the assumption of trend-stationarity for dividends, coupled with this steady state restriction yields:

$$\ln(d_t^*) = \phi_0 + \phi_1 t \tag{7}$$

$$\ln(p_t^*) = \left[\phi_0 + \ln\left(\frac{p^*}{d^*}\right)\right] + \phi_1 t.$$

Thus the variables we work with are logged deviations of prices and dividends from the restricted trajectories given in (7). (See the appendix for details regarding the implementation of this restriction.)

We conclude this section with a characterization of the data. As noted in the introduction, the data are precisely those upon which Shiller (1981) based his analysis of stock-price volatility, updated through 1999. The price series is the S&P's monthly composite stock price index for January, divided by the producer price index (January PPI starting in 1900, annual average PPI from 1871 - 1899); the dividend series represents total dividends for the calendar year accruing to the portfolio represented by the stocks in the index, divided by the average PPI for the year.

As noted, the sample average of the price/dividend ratio is 22.26, implying an annual return of 4.78%; under the CRRA and habit/durability specifications, this implies a corresponding value of β of 0.954. Logged deviations of prices and dividends from the trends estimated subject to (7) for this price/dividend ratio are illustrated in Figure 1. The two series are closely correlated (0.673), and the standard deviation of the price series is nearly twice that of the dividend series (0.406 versus 0.217). In our assessment of the ability of the models to characterize these data, we focus on four summary statistics: the standard deviation of prices and dividends (σ_p, σ_d); the ratio σ_d/σ_p; and the correlation between prices and dividends $(corr(p, d))$. In order to characterize statistical properties of these statistics, we calculated their posterior distributions using a six-lag vector autoregressive representation (VAR) (estimated using the detrended series depicted in Figure 1); this was done using an uninformative prior specified over the VAR parameters. The dashed lines in Figure 3 depict marginal posterior distributions over these statistics; and Table 2 reports posterior means and standard deviations. Note from Figure 3 that these distributions feature considerable skewness, particularly the distribution of σ_p. As a result, e.g., the posterior mean of σ_p (0.62) lies considerably above its sample average of 0.406.

As noted in the introduction, two features of these data constitute the crux of the excess-volatility puzzle: the successful model must account simultaneously for the high volatility of fluctuations in asset prices relative to dividends, and the close correlation observed between movements in these series. We turn now to an assessment of the ability of the alternative preference specifications to account for this behavior.

4 Results

4.1 Estimation

Posterior means and standard deviations obtained for all specifications are presented in Table 1; posterior means and standard deviations of summary statistics corresponding to all models are presented in Table 2. For the CRRA specification, marginal posterior distributions of parameters and summary statistics are graphed in Figures 2 and 3; for the habit/durability specification, these distributions are graphed in Figures 4 and 5; and for the self-control specification, these distributions are graphed in Figures 6 and 7. We begin here with a discussion of the parameter estimates obtained for each model; we then turn to a characterization of fit.

Consider first parameter estimates obtained for the CRRA specification. The posterior distribution of β is centered very close to the prior (with respective means of 0.957 and 0.96), but is much more tightly distributed (its standard deviation of 0.004 is five times less than the prior's). In contrast, the posterior and prior dispersions of γ are similar (respective standard deviations are 0.814 and 1), while the posterior distribution is moderately right-shifted relative to the prior (respective means are 2.884 and 2). Thus the data indicate a somewhat higher degree of risk aversion than was embodied in the prior, but not by a dramatic amount: the means differ by less than one prior standard deviation. Regarding the shocks, the persistence of the endowment shock is greater than that of the dividend shock (posterior means of ρ_d and ρ_e are 0.876 and 0.913, with standard deviations of 0.033 and 0.026), and the shocks are positively correlated (the posterior mean of $corr(\varepsilon_d, \varepsilon_e)$ is 0.313, with posterior standard deviation of 0.124). Regarding the size of the shocks, the posterior distributions of $\sigma_{\varepsilon d}$ and $\sigma_{\varepsilon e}$ have similar means (0.114 and 0.097, respectively); but the standard deviation of $\sigma_{\varepsilon e}$ is five times that of $\sigma_{\varepsilon d}$, and is skewed substantially to the right. Finally, as noted above, the prior and posterior distributions of $\eta = e^*/d^*$ are virtually indiscernible, indicating that the data are uninformative regarding the location of this ratio. This calls into question the general sensitivity of our results to this prior specification, an issue we return to below.

To gain intuition for these results, it is important to keep in mind that the empirical problem confronting the model is largely two-dimensional: simultaneously account for the relatively high volatility of stock prices, and the close correlation observed between prices and dividends. Four parameters play a critical role in confronting this problem. First, large values of $\sigma_{\varepsilon e}$ and η are useful in helping the model along the former dimension: by respectively increasing the volatility and importance of the household's endowment in the budget constraint, these parameters serve to increase the volatility of stock prices without requiring an increase in the volatility of dividends. But these effects are harmful along the latter dimension: they serve to weaken the correlation between movements in prices and dividends. However, this harm can be offset given a large corresponding value of $corr(\varepsilon_d, \varepsilon_e)$, since this of course boosts the correlation between dividends and the endowment process, and thus between dividends and prices. Finally, large values of γ provide help along the former dimension by boosting the volatility of the household's marginal rate of substitution for a given volatility of consumption. But this again serves to weaken the correlation between prices and dividends, because dividend fluctuations do not constitute the exclusive source of fluctuations in consumption.

These considerations are helpful in understanding, for example, why the posterior distribution of γ is only moderately right-shifted relative to the prior, and why the data clearly favor a positive value for $corr(\varepsilon_d, \varepsilon_e)$. They are also helpful in interpreting the posterior correlations estimated between the parameters of the model. Among the eight parameters featured in this version of the model, there are four nontrivial posterior correlations (reported in Table 1): between γ and $\sigma_{\varepsilon e}$ (-0.836); γ and $corr(\varepsilon_d, \varepsilon_e)$ (-0.302); $\sigma_{\varepsilon e}$ and $corr(\varepsilon_d, \varepsilon_e)$ (0.165); and $corr(\varepsilon_d, \varepsilon_e)$ and η (0.520). So for example, given a relatively large value of $\sigma_{\varepsilon e}$, $corr(\varepsilon_d, \varepsilon_e)$ is also relatively large, and γ is relatively small: adjustments in these latter parameters help combat the decreased dividend/price correlation associated with the large value of $\sigma_{\varepsilon e}$.

In light of this correlation pattern, it is easy to convey the implications of adjusting the prior specified over η. Decreasing its prior mean/standard deviation from 10/5 to 5/3 leads to an approximate one-standard-deviation decrease in the posterior mean of $corr(\varepsilon_d, \varepsilon_e)$. Also, the posterior mean of $\sigma_{\varepsilon e}$ rises slightly (by approximately a half-standard-deviation), which apparently serves to offset the volatility reduction in prices corresponding to a decrease in η. Posterior means of the remaining parameters are affected inappreciably by this change.

Consider now the estimates obtained using the habit/durability specification. First, the estimates obtained for the parameters common to both specifications are quite similar. In fact, the posterior mean of ρ_e lies one posterior standard deviation below the mean obtained using CRRA preferences (at 0.889), but the remaining means are all within one posterior standard deviation of their CRRA counterparts. Anticipating the discussion of fit to follow, two additional differences in parameter estimates are noteworthy: the posterior mean of $\sigma_{\varepsilon e}$ falls by nearly one standard deviation (to 0.079), and the posterior mean of $corr(\varepsilon_d, \varepsilon_e)$ rises by nearly one standard deviation (to 0.414). All three of these changes work to generate less-volatile price fluctuations, and greater correlation between dividends and prices.

Regarding the habit/durability parameters, our results are not supportive of the presence of habit: the posterior mode of α is zero (as is that of θ). However, there is clear support for durability. While the posterior of δ is shifted towards zero relative to the prior, the shift is slight: posterior and prior means are 0.0532 and 0.0625, respectively. Also, the height of the posterior at its mode (approximately 0.06) is roughly seven times higher that at zero. The implication of these estimates is that the data are distinctly supportive of a mild degree of local consumption substitutability in the context of this model. Using alternative specifications of preferences, Dunn and Singleton (1986), Gallant and Tauchen (1989), and Eichenbaum and Hansen (1990) also find evidence of local substitutability using monthly aggregate consumption and return data.

Our results of course do not constitute a "rejection" of the notion that consumption is in fact habit-forming. Rather, they indicate that in the context of this environment, the model's characterization of the behavior of the aggregate stock price and dividend data upon which we condition in obtaining these estimates is not enhanced by the generalization of preferences along the habit dimension. Examples of contrasting results obtained using different environments, estimation procedures, and data sets are plentiful. For example, using the same full-information approach to estimation we employ, Otrok (2001) obtained posterior evidence supportive of the presence of both habit and durability in an RBC environment estimated using quarterly observations on non-durable consumption and services and non-residential investment.

Heaton's (1985) analysis of habit/durability is the most closely related to this aspect of our analysis, thus a careful comparison of our results with his is warranted. Heaton studied an asset-pricing environment very similar to ours: it too featured a representative household, an exogenous endowment process, and the same preference specification we consider (indeed, our specification is adopted directly from his). Heaton's results were obtained using a two-step SMM procedure. In the first step, he estimated an autoregressive process for monthly observations on the growth rates of non-durable consumption and services and an annualized dividend series corresponding with CRSP value-weighted returns. In the second step he conditioned on this estimated driving process, and obtained estimates of the same set of preference parameters we consider, using as an objective function variances, cross-correlations, and auto-correlations of monthly CRSP value-weighted returns and returns on one-month T Bills.

Heaton obtained non-zero point estimates for each of the durability and habit parameters, thus at first blush it appears that our results are at odds with his. However, as he carefully notes, it is difficult to make a sharp assessment of the statistical significance of his estimates. For example, his point estimate of α is 0.672, with a corresponding standard error of 1.268. While this is suggestive of statistical insignificance, the fact that α is truncated from above at one raises the possibility that the standard error may provide a poor approximation of the precision of this estimate. Similar observations pertain to his estimates of δ and θ. This problem motivates an analysis of the impact on the fit statistic (sample size times SMM criterion function) of imposing zero restrictions on the preference parameters. While even this is problematic (since this characterization of fit does not correspond with a known asymptotic distribution), it reveals a pattern of results similar to ours. Specifically, relative to the unrestricted model, the imposition of a zero restriction on δ generates a 180% increase in the fit statistic, while the imposition of a zero restriction on α generates only a 29% increase. Thus like us, Heaton finds relatively strong support for durability in his analysis, but weaker support for habit persistence. And as we shall see when we discuss fit below, our results are similar to Heaton's in a second regard: in neither case do we find that habit/durability resolves the excess volatility puzzle. (Our failure to find a significant role for habit formation in this setting obtains despite the fact that we work with annual data: as noted, e.g., by Ferson and Constantinides (1991), the use of annual data is helpful in detecting slowly developing habit formation.)

Consider now the estimates obtained using the self-control specification. In this case, the estimates obtained for the parameters common to both the self-control and CRRA specifications are relatively distinct. Most notably, the posterior mean of β increases by approximately one standard deviation under the self-control specification; the means of γ and η decrease by approximately one standard deviation; and the mean of $\sigma_{\varepsilon e}$ increases by more than 1.5 standard deviations. Again anticipating the discussion of fit to follow, the decreases in γ and η will serve to increase the correlation between prices and dividends in the model, while dampening the volatility of prices; the increase in $\sigma_{\varepsilon e}$ will have the opposite effect. This raises the question of why these offsetting factors arise. The key lies in reconsidering the relationship between steady state prices and dividends given in (6). Movements away from zero by the self-control parameters λ and ϕ cause sharp declines in the steady state price/dividend ratio; decreases in γ and η, and increases in β, can serve to offset these effects. Thus it appears that the altered estimates of γ, η and β arise from an empirical preference for non-zero values of λ and ϕ; the posterior estimates of these self-control parameters confirm this observation.

The posterior mean of ϕ is 0.3, roughly one posterior standard deviation below its prior mean, but still significantly higher than zero. And the posterior mean of λ is 0.00286, which coincides closely with its mode, and lies approximately 1.5 posterior standard deviations above the prior mode of zero. To appreciate the quantitative significance of this estimate, consider two alternative calculations. First, returning to (6), if we fix γ, η, β and ϕ at their posterior means, an increase of λ from zero to 0.00286 implies a decrease in the steady state price/dividend ratio from approximately 28 to 20.46, implying an increase in annual average returns from 3.57% to 4.89%. In other words, a household endowed with the self-control preferences associated with these parameter estimates would demand an annual "return premium" over a household endowed with CRRA preferences of 1.32%, representing a 36% difference. Second, consider the quantity of steady state consumption a household endowed with the temptation specification we estimate would be willing to sacrifice in order to be rid of this temptation. This is the value x such that

$$u(\overline{c} - x) = u(\overline{c}) + v(\overline{c}) - v(\widetilde{c}),$$

where here $\widetilde{c}$ denotes the steady state value of temptation consumption. Using posterior means of our parameter estimates, the implied value of x amounts to 5.25% of $\overline{c}$. Thus we obtain support for the presence of a quantitatively significant temptation effect in the data.

4.2 Fit

We now characterize implications of these parameter estimates for the behavior of the variables included in the models. As noted, we summarize this behavior with the following set of statistics: standard deviations of prices, dividends, and endowments (σ_p, etc.); the ratios ($\sigma_d/\sigma_p, \sigma_e/\sigma_p$); and cross-correlations between these variables. To evaluate fit, we compare posterior distributions of the statistics associated with the observable variables ($\sigma_p, \sigma_d, \sigma_d/\sigma_p, corr(p,d)$) with their empirical counterparts (the posteriors obtained using the flat-prior VAR). The posterior distributions of these statistics are summarized in Table 2, and graphed in Figures 3, 5, and 7 (CRRA preferences, habit/durability preferences, self-control preferences).

Visual comparisons of model and VAR posteriors provide one way to evaluate fit: loosely, the greater the correspondence, the better the fit. Beyond this, we rely on two formal measures of correspondence: the confidence-interval-criterion (CIC) measure proposed by DeJong, Ingram and Whiteman (1996); and the posterior-odds measure proposed by Geweke (1999). Each measure enables head-to-head fit comparisons for sets of alternative models, including non-nested ones.

Regarding the CIC measure, let s be a summary statistic that can be calculated both from the model and the VAR; and let $P(s)$ and $V(s)$ denote the marginal posterior distributions over s associated with the model and VAR. Finally, let [a b] denote the inter-κ quartile range of $V(s)$ (i.e., a is the $\kappa/2$ quantile of $V(s)$, and b is the $1-\kappa/2$ quantile of $V(s)$). Then the CIC measure associated with s we employ is given by

$$CIC(s) = \int_a^b P(s)ds;$$

thus $CIC(s)$ measures the proportion of the distribution of $P(s)$ that lies within the inter-κ quartile range of its empirical counterpart $V(s)$.

Regarding Geweke's (1999) posterior odds measure, let m denote a vector of summary statistics, $PA(m)$ the posterior distribution of m associated with model A, and $V(m)$ the posterior distribution of m associated with the VAR. Then assigning even prior odds to models A and B, the posterior odds in favor of model A relative to B are given by

$$PO_{A,B} = \frac{\int P_A(m)V(m)dm}{\int P_B(m)V(m)dm};$$

thus $PO_{A,B}$ provides a characterization of the relative overlap between the VAR density over m

and the posterior densities associated with models A and B. The greater the overlap exhibited by model A relative to model B, the greater will be $PO_{A,B}$.

Evaluation of the integrals required to calculate $PO_{A,B}$ is facilitated by kernel density approximation. Let $m_A(i)$ denote the *ith* of M drawings of m obtained from $P_A(m)$ (or in our case, a suitably weighted drawing from the importance density associated with $P_A(m)$), and $m_V(j)$ denote the *jth* of N drawings obtained from $V(m)$. Then the numerator of $PO_{A,B}$ may be approximated using

$$\frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}K(m_A(i), m_V(j)),$$

and likewise for the denominator. Here, we use a Normal density kernel with window width chosen to minimize approximate mean integrated square error, following Silverman (1986, equation 3.31).

A key difference between the CIC and PO measures is that the CIC measure facilitates comparisons of marginal densities of a single statistic in isolation, and thus is useful in highlighting performance along a specific dimension. In contrast, the PO measure enables comparisons of the joint behavior of the entire vector of summary statistics. Thus if there is significant interaction between the statistics, this will be picked up by the PO measure but not the CIC measure, and inferences based on the two measures can potentially vary.

Turning to our results, note first that all three models provide virtually identical characterizations of the volatility of dividends, and that the characterizations correspond very closely with the data: e.g., $CIC(\sigma_d)$ is close to 1 in all cases. Differences across models are evident, however, in the remaining series.

Consider first a comparison of the CRRA and habit/durability specifications. Note that the volatility of the endowment is considerably lower in the habit/durability specification: respective posterior means (standard deviations) of σ_e are 0.175 (0.061) and 0.251 (0.116); moreover, the posterior of σ_e obtained using the CRRA specification is distinctly right-skewed. In light of the lower estimates of the parameters $(\rho_e, \sigma_{\varepsilon e})$ obtained for the habit/durability specification, this difference is not surprising.

The implication of this behavior for prices is also not surprising: prices are far smoother in the habit/durability specification than in the CRRA specification, and more closely correlated with dividends. As a result, the CRRA specification does a relatively good job in characterizing price volatility, at the cost of matching the data along the correlation margin; and the opposite is true of the habit/durability specification. For example, $CIC(\sigma_p)$ is 0.616 in the CRRA specification, but only 0.214 in the habit/durability specification; while $CIC(corr(p,d))$ is 0.865 in the CRRA specification, and 0.998 in the habit/durability specification. Thus like Heaton (1985), we find that habit/durability preferences are not capable of accounting for stock-price volatility.

This is reflected further in our posterior-odds comparison: the odds favoring the CRRA model relative to the habit/durability model are nearly 2:1 (1.987). Given the distinct evidence favoring a non-zero value of the durability parameter δ provided by its marginal posterior distribution illustrated in Figure 4, this result is surprising at first blush. However, it is important to interpret this result appropriately: it implies that the CRRA specification provides a superior characterization of the four summary statistics we have chosen to highlight on a weighted-average basis relative to the habit/durability specification. One explanation for this could be, for example, that the likelihood function places greater emphasis on other features of the data than on the summary statistics we highlight. Nevertheless, it is clear that the habit/durability extension does not enhance the model's characterization of stock-price volatility in this environment.

The CRRA and self-control specifications are more closely comparable. Once again, an obvious difference lies along the σ_e dimension, but in this case the self-control specification delivers a relatively large value for this statistic: 0.391 (0.129) versus 0.251 (0.116). This is the result of the relatively large estimate of $\sigma_{\varepsilon e}$ obtained using the self-control specification. Recall that the increased estimate of $\sigma_{\varepsilon e}$ arose under the self-control specification from a need to offset the impact on predicted price volatility of the decreased estimates of γ and η.

Regarding fit, the self-control specification marginally outperforms the CRRA specification according to the CIC criterion along both the volatility and correlation dimensions: $CIC(\sigma_p)$ is 0.624 in the self-control specification (compared with 0.616), and $CIC(corr(p,d))$ is 0.902 in the self-control specification (compared with 0.865). The posterior-odds comparison yields a more dramatic difference: the odds favoring the self-control over the CRRA specification are approximately 2.4:1.

Thus it appears that the self-control specification does a relatively good job in characterizing the joint behavior of the statistics we use to evaluate fit.

Taken as a whole, the parameter estimates and evaluations of fit we obtain are supportive of a positive role for self-control preferences in helping account for the interaction of aggregate stock prices and dividends. The data support the presence of a significant temptation effect according to our posterior estimates. And in the presence of temptation, lower values of the coefficient of relative risk aversion γ are needed to characterize volatility; the importance of the exogenous endowment process is also diminished. These results are useful in helping the model to deliver an adequate characterization of the strong correlation pattern observed between stock prices and dividends, while at the same time attempting to characterize stock price volatility. In the context of the simple environment we study, this characterization is not completely successful; nevertheless, it seems to represent a step in the right direction.

5 Conclusion

We have used a simple asset-pricing environment to study the ability of self-control preferences to account for the interaction of aggregate annual stock prices and dividends. Using as benchmarks results obtained using CRRA and habit/durability preferences, we found empirical support for a positive role for temptation in characterizing the behavior of these data. Moreover, parameter estimates obtained using a full-information procedure indicate the presence a quantitatively significant temptation effect. While self-control preferences do not fully account for the level of price volatility observed in the data, our results are nevertheless encouraging, particularly in light of the simple environment upon which our analysis was based.

A Appendix

Here we provide a sketch of the technical details of our estimation procedure. Let μ denote the vector of parameters associated with a given model, and x_t denote the (n x 1) vector of variables in the model, written as logged deviations from steady state values. For a given specification of μ, log-linear approximation of the model yields a first-order system of the form

$$x_t = Fx_{t-1} + G\varepsilon_t \tag{A1}$$

where ε_t denotes the (m x 1) vector of disturbances in the model with a corresponding covariance matrix Σ, and the elements of (F, G, Σ) are functions of μ. Although the system is of dimension n, it is stochastically singular because it includes only m random shocks. Thus the model carries non-trivial predictions for an (m x 1) vector of variables X_t, with a mapping from x_t given by

$$X_t = H^{/} x_t \tag{A2}$$

Coupled with the assumption of normality for ε_t, (A1) and (A2) yield a likelihood function $L(X|\mu)$ that can be evaluated using the Kalman filter. Finally, the specification of a prior distribution $p(\mu)$ yields a posterior distribution for μ via Bayes' Rule:

$$P(\mu|X) \propto L(X|\mu)p(\mu). \tag{A3}$$

The expected value of a general function of interest $g(\mu)$ under the posterior is given by

$$E[g(\mu)] = \frac{\int g(\mu)P(\mu|X)d\mu}{\int P(\mu|X)d\mu}. \tag{A4}$$

(Here, our interest is in the parameters themselves, as well as in various moments of x_t.) In general, the integrals in (A4) must be approximated using numerical integration techniques. Ideally, this is done by generating an artificial sample $\{\mu_k\}$ for k = 1,...,D directly from the posterior density (A3), and approximating (A4) by calculating the average value of $g(\mu)$ obtained using these drawings. But since the likelihood function in this case is in the form of an observer system, it is not possible to generate parameter drawings from its associated posterior distribution. Instead, we proceed via Importance Sampling, which involves generating an artificial sample from a different distribution from which it is possible to sample, and assigning weights to the elements of the sample so that they can be thought of as originating from the posterior distribution of interest. The distribution $I(\mu)$ used to obtain drawings of m is known as the importance density. Given an artificial sample, (A4) is approximated by calculating the weighted average

$$\overline{g}_D = \frac{\sum_{i=1}^{D} g(\mu_i)w(\mu_i)}{\sum_{i=1}^{D} w(\mu_i)}, \tag{A5}$$

where the weight function $w(\mu_i) = P(\mu_i|X)/I(\mu_i)$; $I(\mu_i)$ appears in the denominator of $w(\mu_i)$ to offset the direct influence that $I(\mu)$ has in obtaining the particular drawing μ_i. Given that the support of $I(\mu)$ includes that of $P(\mu|X)$, $\overline{g}_D$ converges almost surely to $E[g(\mu)]$, so long as $E[g(\mu)]$

exists and is finite.

To obtain the results reported in the paper, we specified a multivariate t distribution for $I(\mu)$ for each model to insure that their supports included those of $P(\mu|X)$. Means and covariance matrices for $I(\mu)$ were chosen sequentially. In calculating first-pass approximations of posterior means and covariance matrices of μ, only very few drawings from $I(\mu)$ typically received appreciable posterior weight. Thus after obtaining an initial round of draws, means and covariance matrices of $I(\mu)$ were relocated at the first-pass approximations, and a second round of drawings was obtained. After several rounds moment calculations converged (subject to numerical sampling error) to those used in deriving the results presented in the paper. Of these drawings, that which was assigned the greatest weight received less than three percent of the total assigned weight under the habit/durability specification, and less than one percent under the CRRA and self-control specifications; hence we are confident that our results closely approximate the actual posterior calculations we seek.

Recall that the data we work with are logged deviations of dividends and prices from the restricted trend specification given in (7). Since the restrictions are a function of the parameters μ, the raw data must be re-transformed for every parameter drawing we obtain. Specifically, given a particular drawing of μ, which implies a particular value for p^*/d^*, we regress logged prices and dividends on a constant and linear time trend, imposing the parameter restrictions indicated in (7); deviations of the logged variables from their respective trend specifications constitute the values of X_t used to evaluate the likelihood function for that specific drawing of μ.

We conclude with a note regarding the pricing equation (2b) obtained under the habit/durability specification. This equation relates prices to an infinite sum of leads and lags of consumption, which must be truncated prior to empirical evaluation. Experimentation with alternative truncation points indicates that consumption values beyond four leads/lags receive little weight in this equation under plausible parameterizations, thus our model estimates were obtained using this approximation.

References

[1] DeJong, D.N. [1992]: "Co-Integration and Trend-Stationarity in Macroeconomic Time Series," Journal of Econometrics, Vol. 52: 347-370.

[2] DeJong, D.N., B.F. Ingram and C.H. Whiteman [1996]: "A Bayesian Approach to Calibration," Journal of Business and Economic Statistics, Vol. 14: 1:9.

[3] DeJong, D.N., B.F. Ingram and C.H. Whiteman [2000]: "A Bayesian Approach to Dynamic Macroeconomics," Journal of Econometrics, Vol. 98: 203-223.

[4] DeJong, D.N. and C.H. Whiteman [1991]: "The Temporal Stability of Dividends and Stock Prices: Evidence from the Likelihood Function," American Economic Review, Vol. 81, No. 3, June: 600-617.

[5] DeJong, D.N. and C.H. Whiteman [1994]: "Modeling Stock Prices Without Knowing How to Induce Stationarity," Econometric Theory, Vol. 10: 701-719.

[6] Dunn, K.B. and K.J. Singleton [1986]: "Modeling the Term Structure of Interest Rates Under Nonseparable Utility an Durability of Goods," Journal of Empirical Finance, Vol. 17: 27-55.

[7] Eichenbaum, M.S. and L.P. Hansen [1990]: "Estimating Models with Intertemporal Substitution Using Aggregate Time Series Data," Journal of Business and Economic Statistics, Vol. 8: 53-69.

[8] Ferson, W.E. and G.M. Constantinides [1991]: "Habit Formation and Durability in Aggregate Consumption: Empirical Tests," Journal of Financial Economics, Vol. 29: 199-240.

[9] Gallant, A.R. and G. Tauchen [1989]: "Seminonparametric Estimation of Conditionally Heterogeneous Processes: Asset Pricing and Applications," Econometrica, Vol. 57: 1091-1120.

[10] Geweke, J. [1999]: "Computational Experiments and Reality," University of Iowa working paper.

[11] Grossman, S.J. and R.J. Shiller [1981]: "The Determinants of the Variability of Stock Market Prices," American Economic Review, Vol. 71: 222-227.

[12] Gul, F. and W. Pesendorfer [2000]: "Self-Control and the Theory of Consumption," Mimeo, Princeton University, October.

[13] Heaton, J. [1985]: "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," Econometrica, Vol. 63, No.3, May: 681-717.

[14] Kocherlakota, N [2001]: "Looking for Evidence of Time-Inconsistent Preferences in Asset Market Data," *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 25, No. 3: 13-24.

[15] Krusell, P., B. Kuruscu and A.A.Smith [2002]: "Time Orientation and Asset Prices," in *Journal of Monetary Economics*, 49, Issue 1.

[16] Lucas, R.J. [1978]: "Asset Prices in an Exchange Economy," Econometrica, Vol. 46: 1429-1445.

[17] LeRoy, S.F. and R.D. Porter [1981]: "Stock Price Volatility: Tests Based on Implied Variance Bounds," Econometrica, Vol. 49: 97-113.

[18] Otrok, C. [2001]: "On measuring the welfare cost of business cycles," in *Journal of Monetary Economics*, 47, Issue 1.

[19] Shiller, R.J. [1981]: "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" American Economic Review, Vol. 71, No.3, June: 421-436.

[20] Silverman, B.W. [1986]: Density Estimation. London: Chapman and Hall.

Table 1. Parameter Estimates

	β	γ	δ	θ	α	λ	ϕ	ρ_d	ρ_e	$\sigma_{\varepsilon d}$	$\sigma_{\varepsilon e}$	corr(ε_d,ε_e)	$\bar{e}/\bar{d}$
Prior Mean	0.960	2.000	0.0625	0.0625	UN	0.0000	0.4	0.900	0.900	UN	UN	UN	10.00
CRRA Mean	0.957	2.884	NA	NA	NA	NA	NA	0.876	0.913	0.114	0.097	0.313	10.68
Hab/Dur Mean	0.959	2.754	0.0532	0.0248	0.0004	NA	NA	0.891	0.889	0.114	0.079	0.414	10.57
Self-Control Mean	0.966	1.973	NA	NA	NA	0.00286	0.30	0.880	0.911	0.115	0.154	0.290	5.95
Prior Std. Dev.	0.020	1.000	0.030	0.030	UN	0.01	0.2	0.050	0.050	UN	UN	UN	5.00
CRRA Std. Dev.	0.004	0.814	NA	NA	NA	NA	NA	0.033	0.026	0.007	0.035	0.124	4.39
Hab/Dur Std. Dev.	0.003	0.697	0.0245	0.0315	0.0036	NA	NA	0.032	0.024	0.007	0.028	0.132	4.586
Self-Con. Std. Dev.	0.007	0.422	NA	NA	NA	0.00182	0.14	0.027	0.023	0.007	0.034	0.108	3.22

Posterior Correlations:	γ, $\sigma_{\varepsilon e}$	γ, corr(ε_d,ε_e)	$\sigma_{\varepsilon e}$, corr(ε_d,ε_e)	corr(ε_d,ε_e), $\bar{e}/\bar{d}$
CRRA Preferences	-0.836	-0.302	0.165	0.520
Hab/Dur Preferences	-0.899	-0.301	0.200	0.609
Self-Control Preferences	-0.718	-0.343	0.047	0.553

Notes: UN denotes “uninformative prior”; NA denotes “not applicable”.

Table 2. Summary Statistics

	σ_p	σ_d	σ_e	σ_d/σ_p	σ_e/σ_p	corr(p,d)	corr(p,e)	corr(d,e)
VAR Mean	0.620	0.268	NA	0.448	NA	0.672	NA	NA
CRRA Mean	0.434	0.244	0.251	0.571	0.568	0.536	0.960	0.301
Hab/Dur Mean	0.366	0.261	0.175	0.710	0.482	0.669	0.926	0.408
Self-Con. Mean	0.418	0.248	0.391	0.608	0.941	0.564	0.939	0.282
VAR Std. Dev.	0.183	0.065	NA	0.095	NA	0.178	NA	NA
CRRA Std. Dev.	0.067	0.041	0.116	0.101	0.212	0.079	0.043	0.120
Hab/Dur Std. Dev.	0.041	0.051	0.061	0.082	0.170	0.071	0.063	0.130
Self-Con. Std. Dev.	0.069	0.035	0.129	0.125	0.296	0.101	0.079	0.106
CRRA CIC	0.616	0.978	NA	0.549	NA	0.865	NA	NA
Hab/Dur CIC	0.214	0.946	NA	0.039	NA	0.998	NA	NA
Self-Con. CIC	0.624	0.975	NA	0.479	NA	0.902	NA	NA

Posterior Odds: CRRA versus Hab/Dur: 2 : 1;   CRRA versus Self-Control 1 : 2.4

Notes: VAR statistics summarize flat-prior posterior distributions associated with a six-lag vector autoregression; CIC denotes the "confidence interval criterion" of DeJong, Ingram and Whiteman (1996); σ_x denotes the standard deviation of the logged deviation of x from its steady state value; and corr(x,y) denotes the correlation observed between logged deviations from steady state values of x and y.

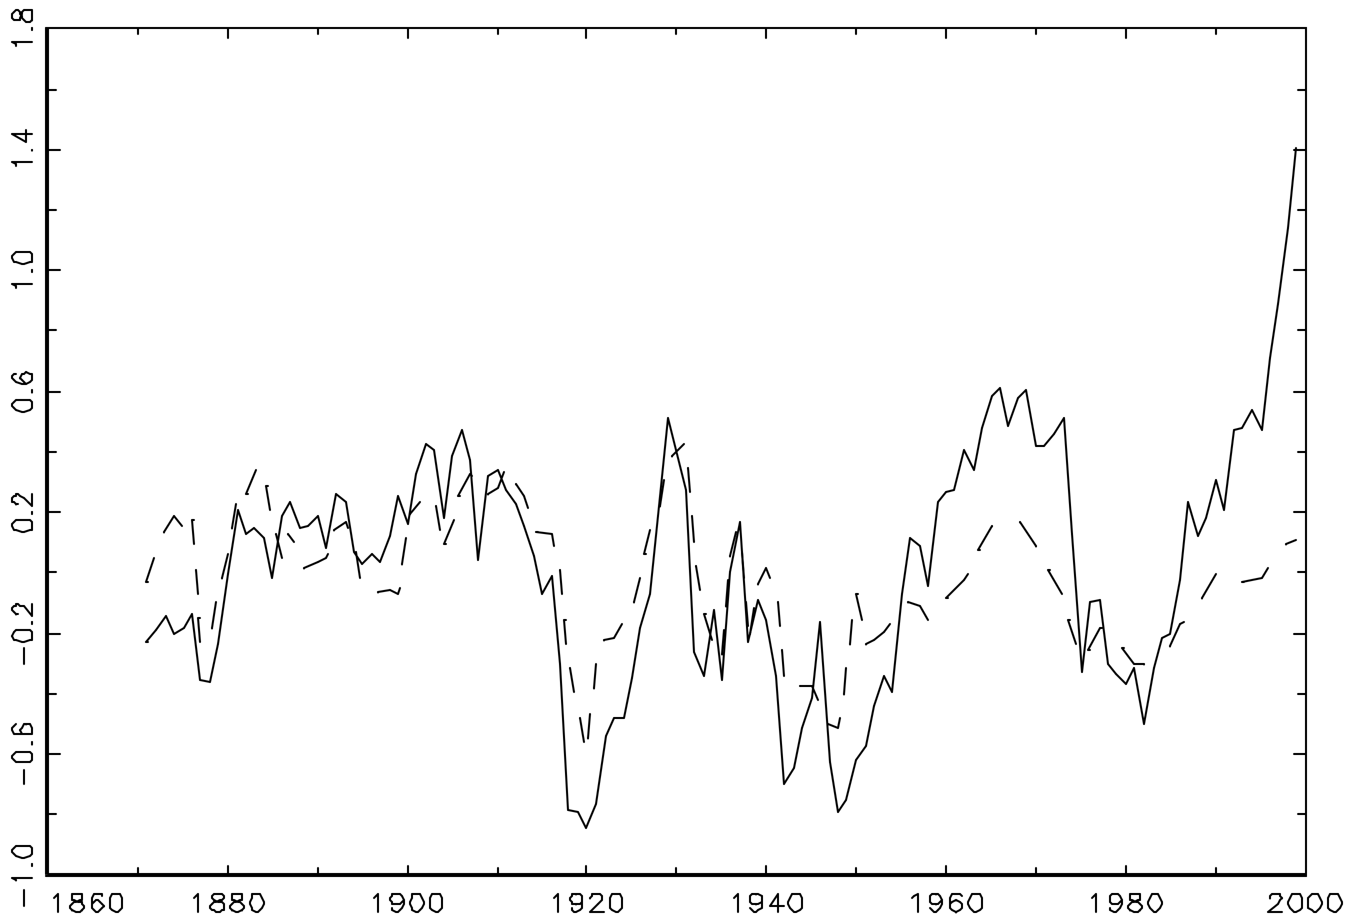

Figure 1. Stock Price and Dividend Data (Logged Deviations from Trend)
Prices: Solid Line Dividends: Dashes

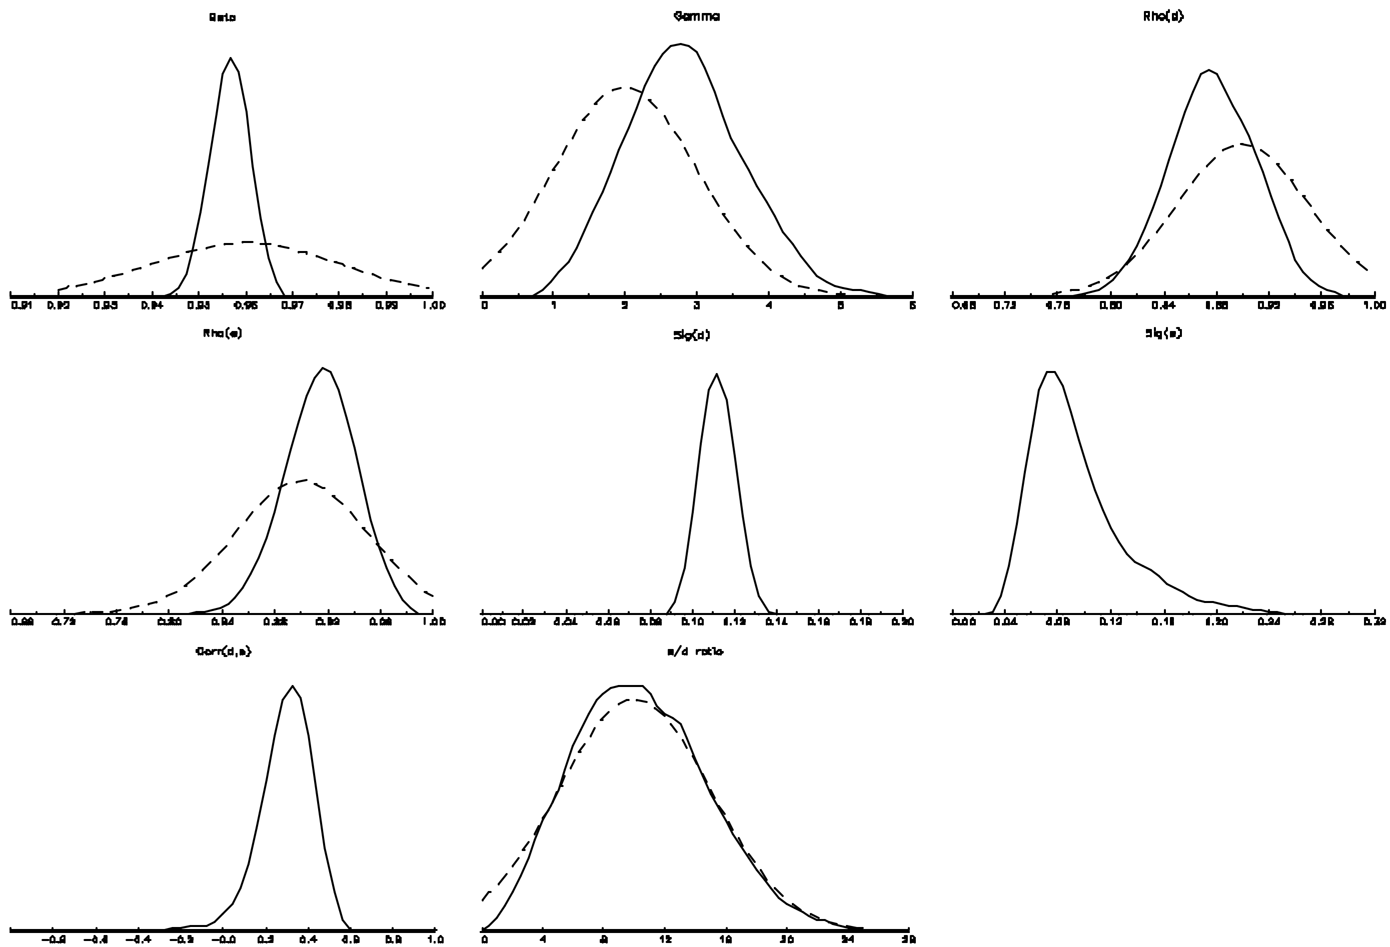

Figure 2. Parameter Estimates, CRRA Preferences

Posteriors: Solid Lines   Priors: Dashes

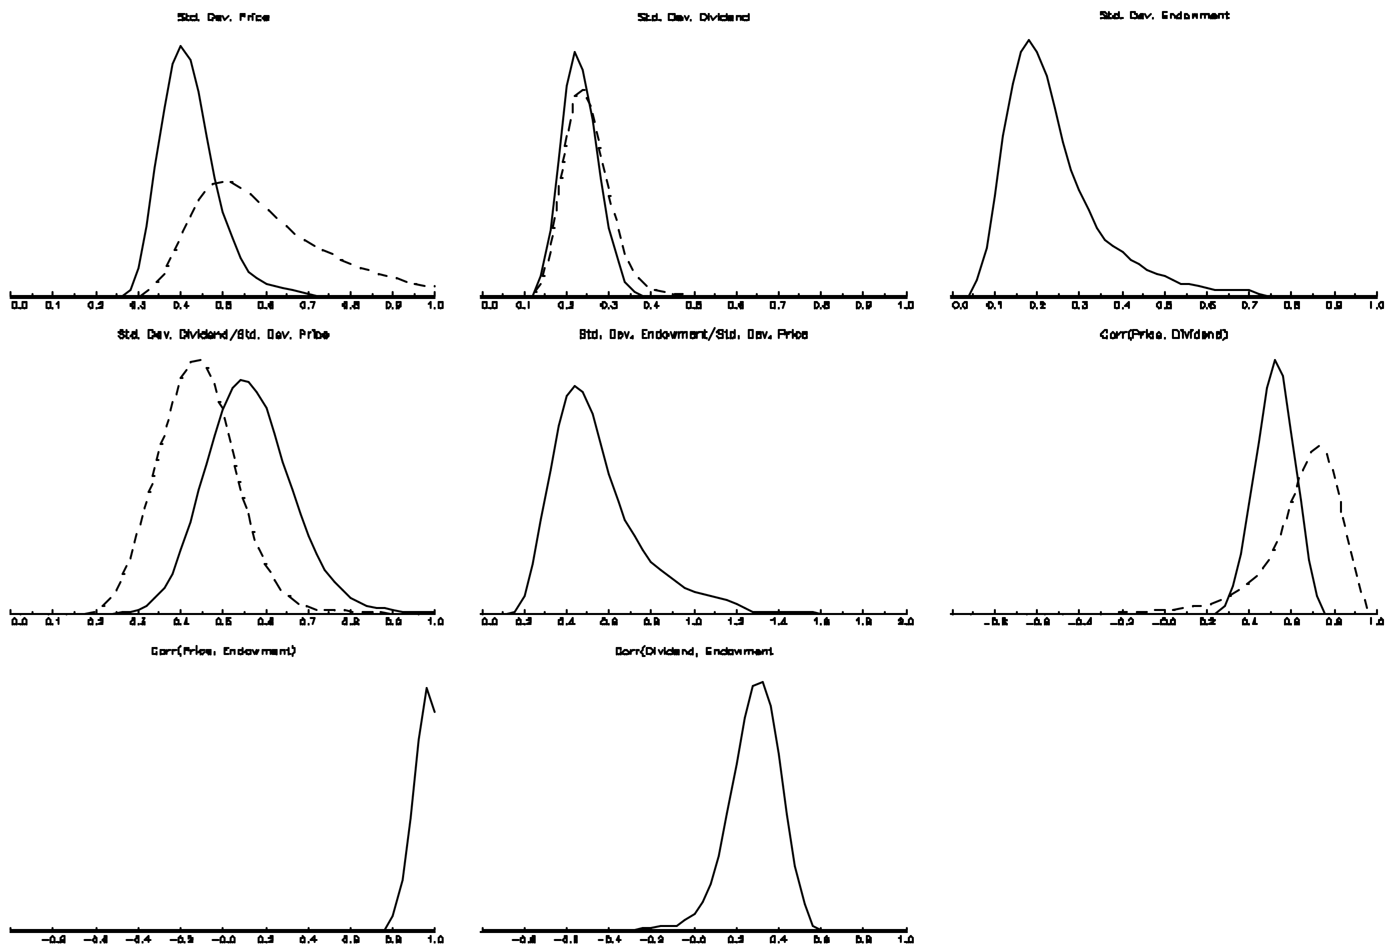

Figure 3. Summary Statistics, CRRA Preferences

Model: Solid Lines VAR: Dashes

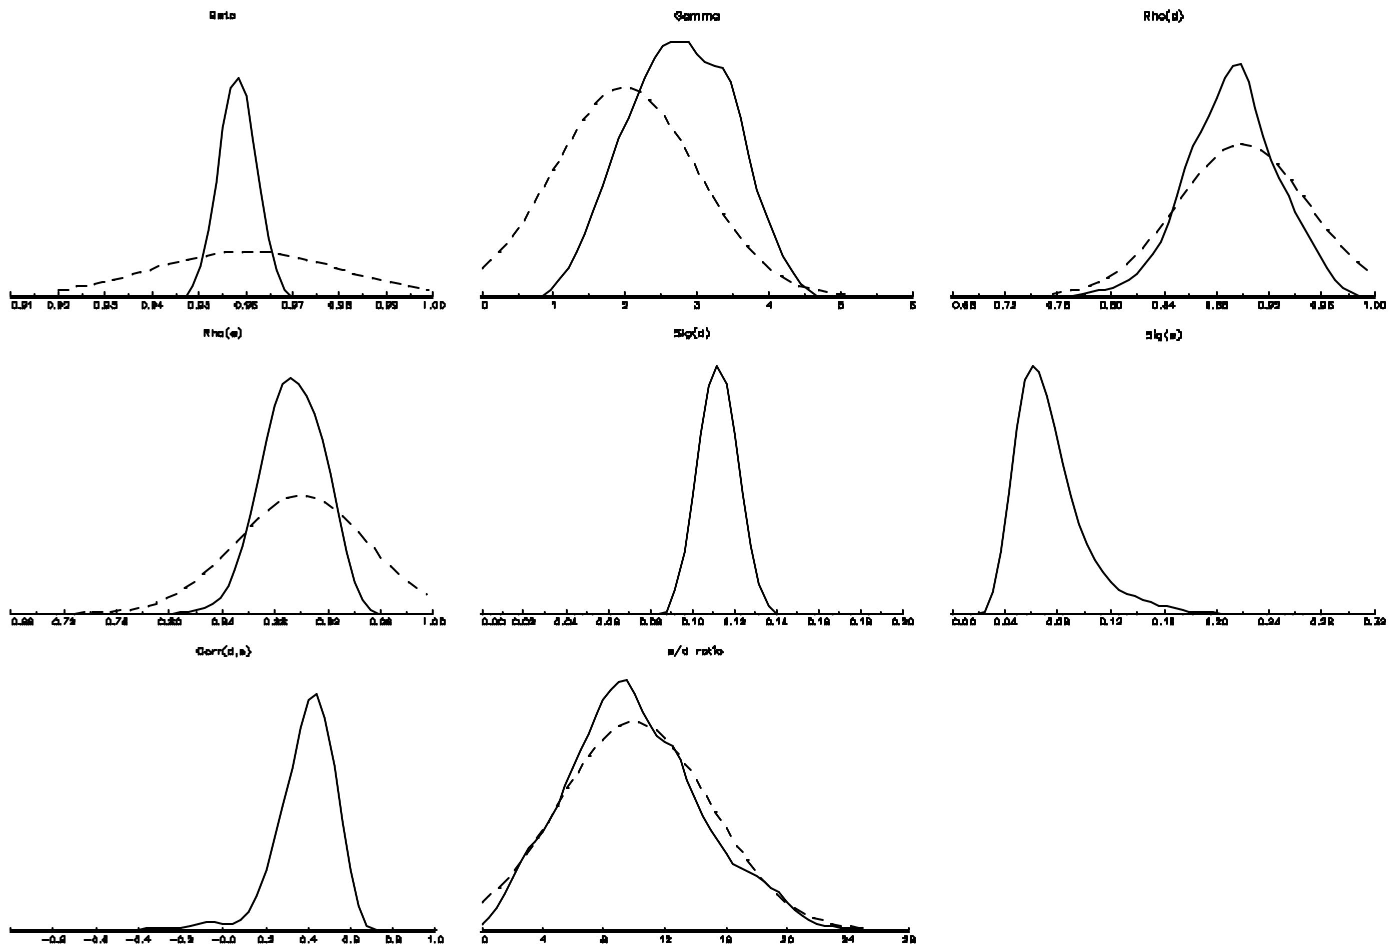

Figure 4. Parameter Estimates, Habit/Durability Preferences

Posteriors: Solid Lines   Priors: Dashes

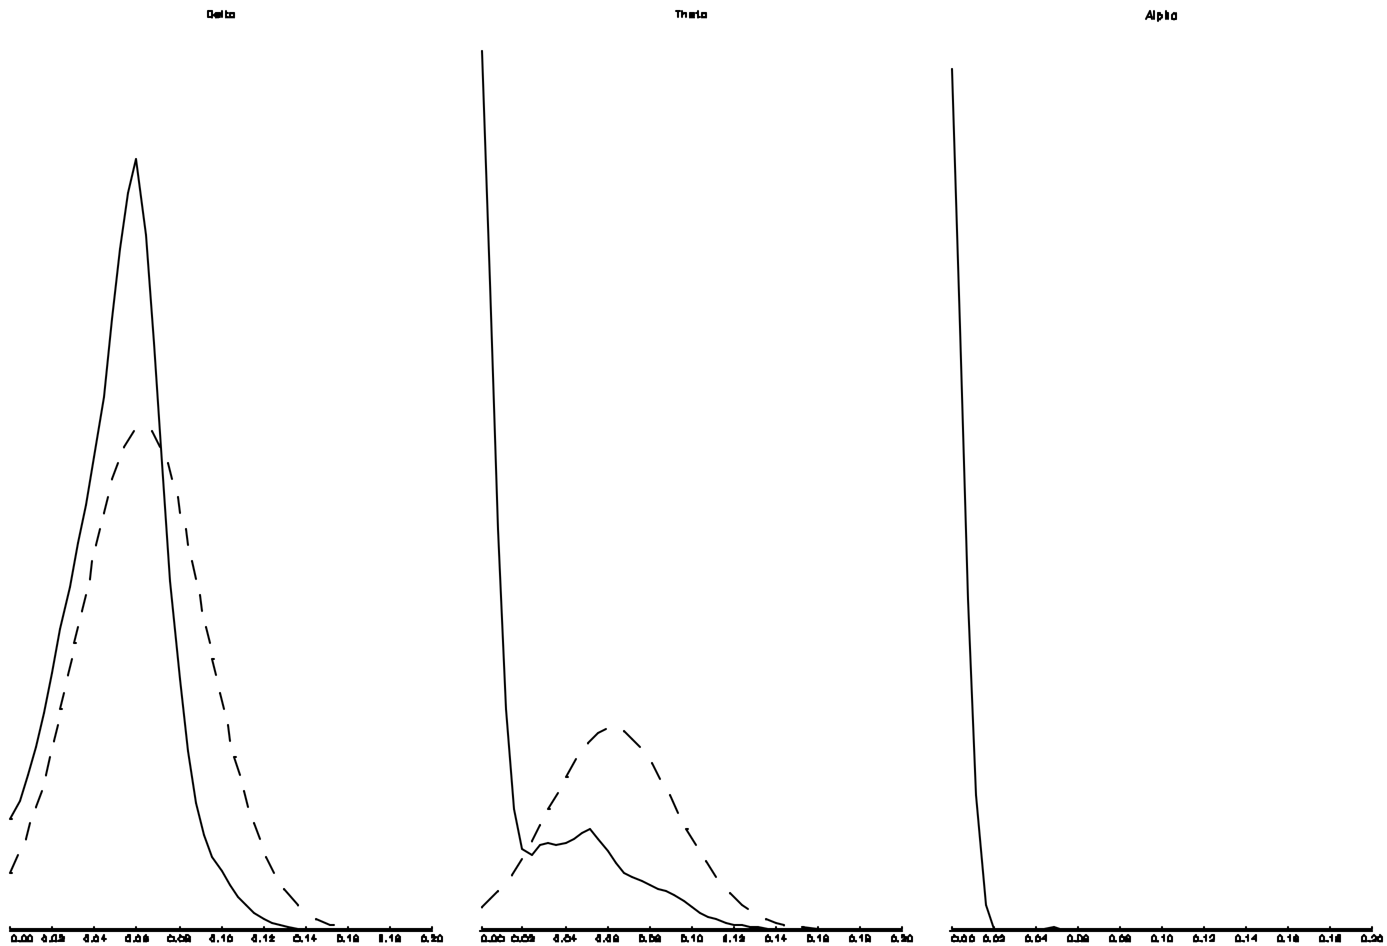

Figure 4, continued

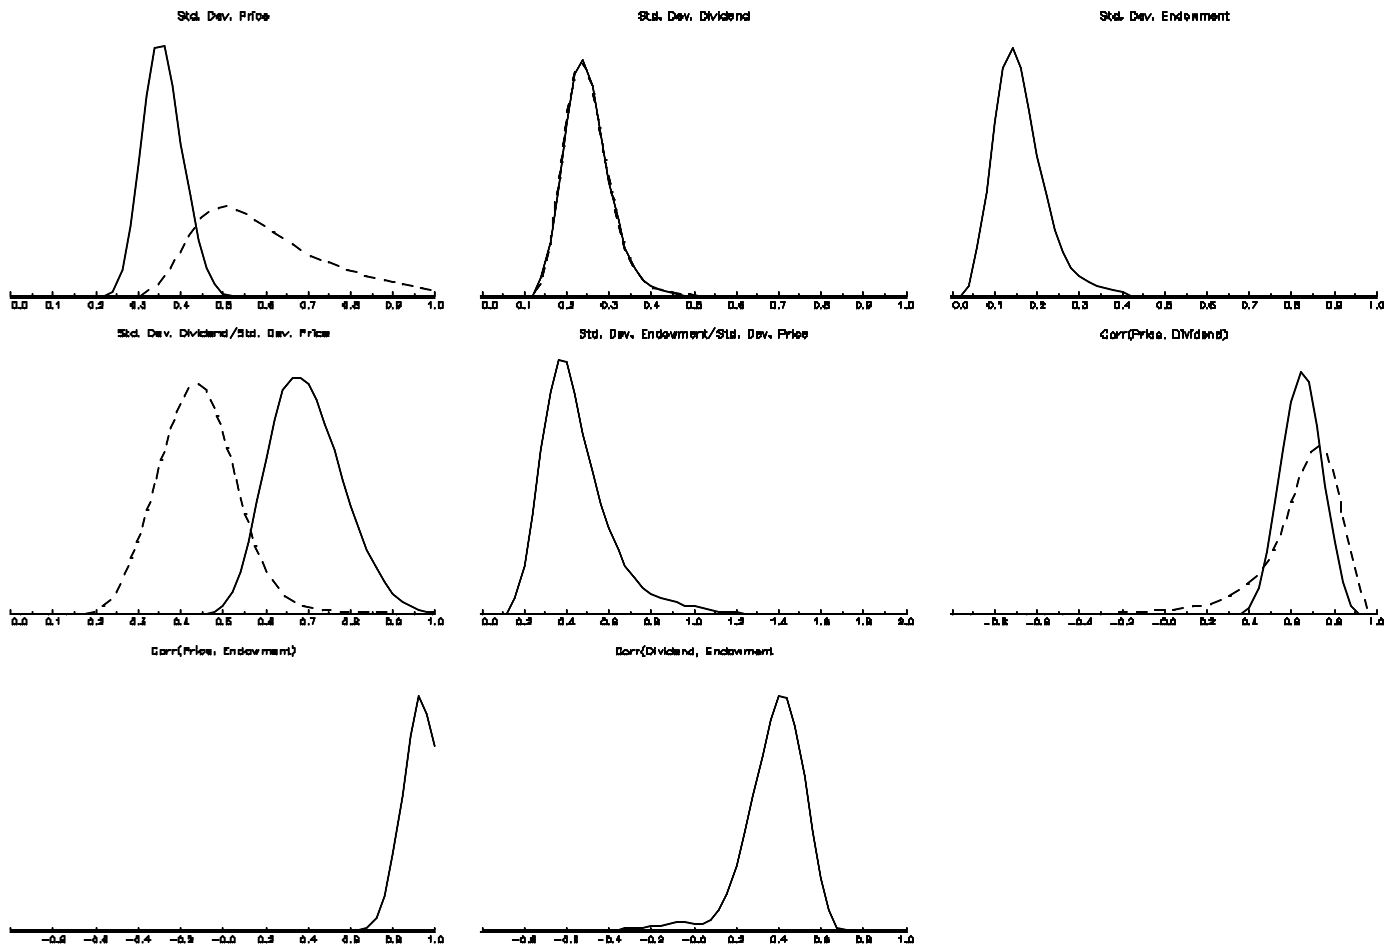

Figure 5. Summary Statistics, Habit/Durability Preferences

Model: Solid Lines VAR: Dashes

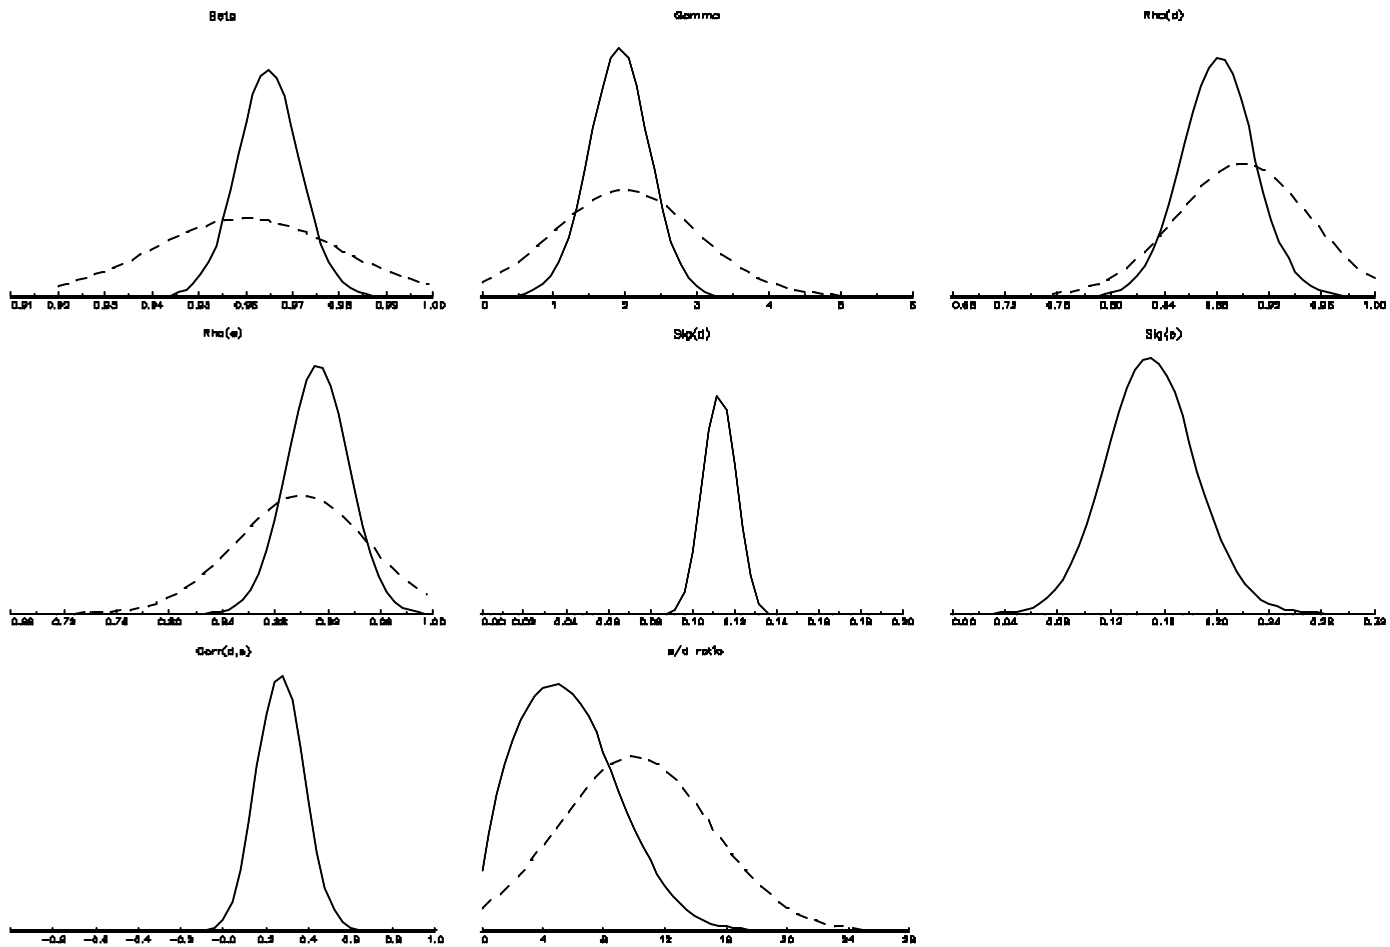

Figure 6. Parameter Estimates, Self-Control Preferences

Posteriors: Solid Lines   Priors: Dashed Lines

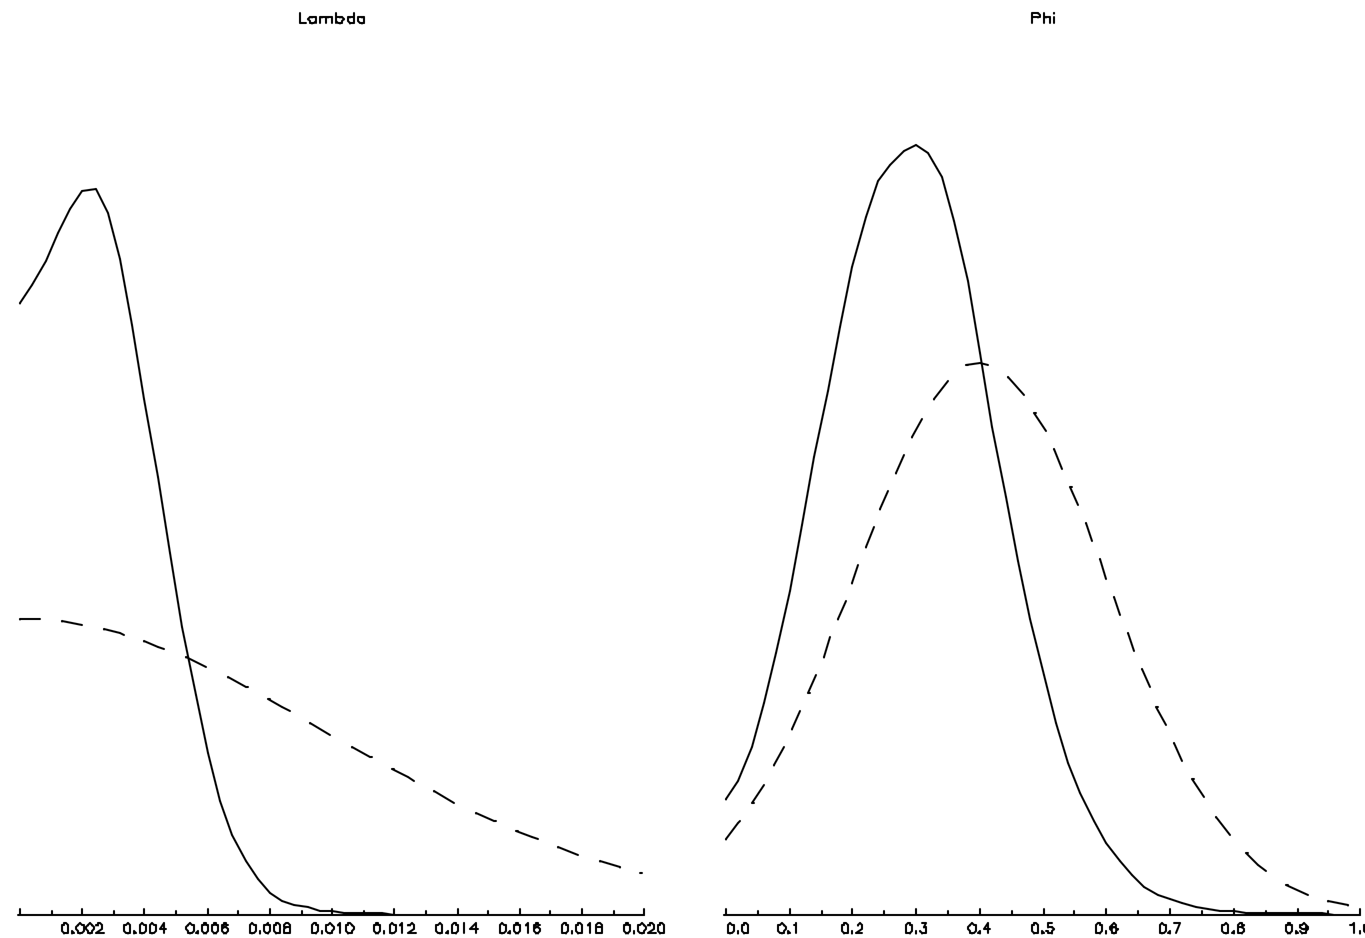

Figure 6, continued

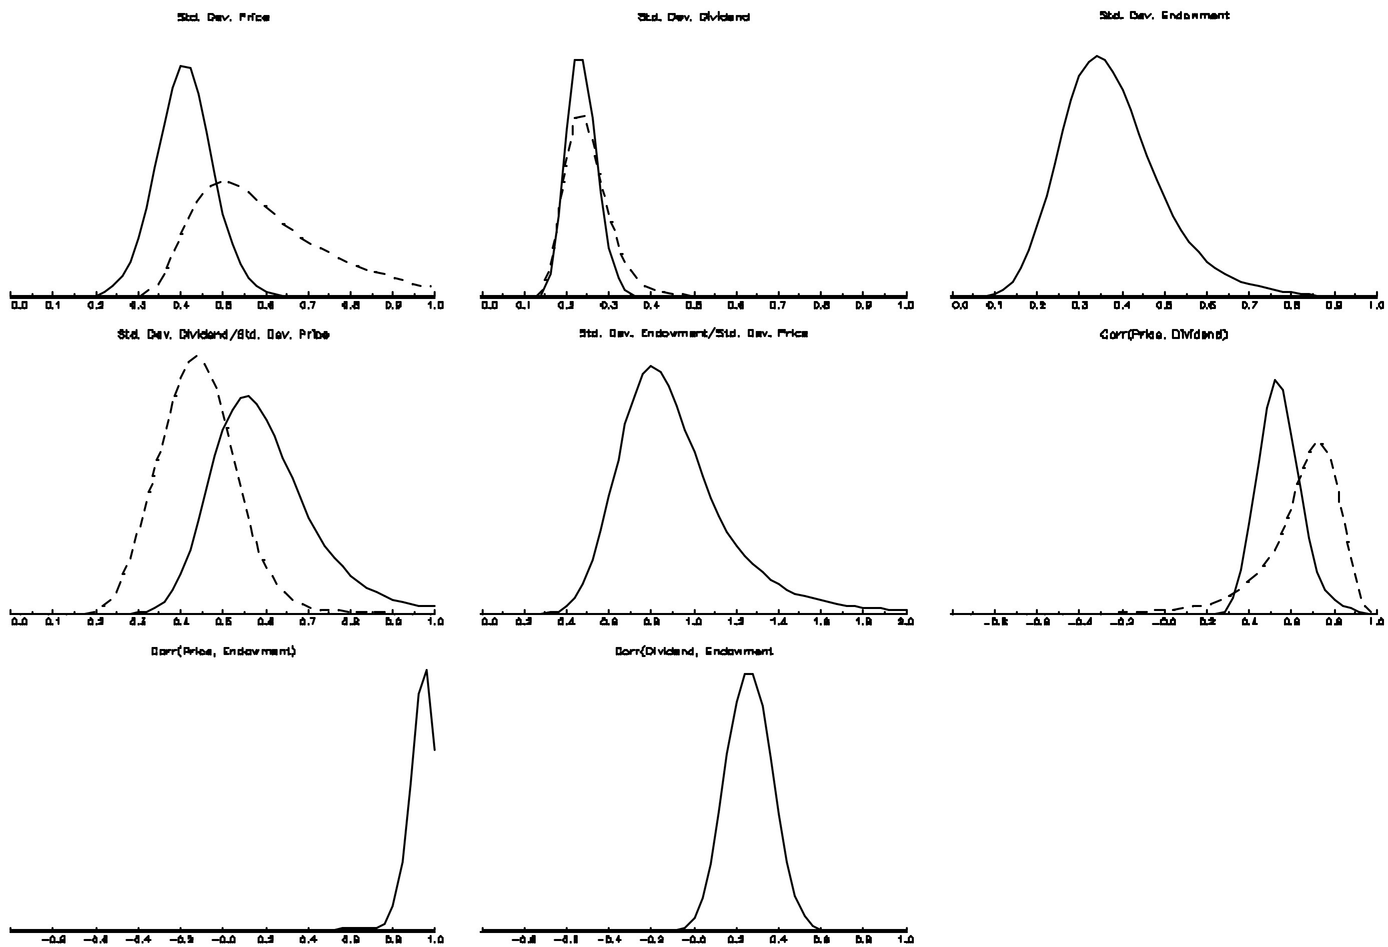

Figure 7. Summary Statistics, Self-Control Preferences

Model: Solid Lines VAR: Dashes