By Force of Habit and Cyclical Leverage*

David Schreindorfer
W.P. Carey School of Business
Arizona State University

October 1, 2023

Abstract

I show that some of Campbell and Cochrane's (1999) main predictions about asset price dynamics result from an inaccuracy in their model's numerical solution. When solved accurately, volatility is a hump-shaped function of the model's state and falls in recessions. The leverage effect counterfactually disappears in recessions. Risk premia are substantially less cyclical and returns less predictable than reported in the original paper. At the 7-year horizon, the predictive R-squared falls from 30% for the original solution to 10% for the accurate solution. However, I show that the original findings can be restored by augmenting the model with countercyclical leverage.

JEL Classification: G12, G13, G33

Keywords: External habits, countercyclical volatility, return predictability, leverage, numerical solution methods.

*Contact: `david.schreindorfer@asu.edu`. I am grateful to John Campbell and John Cochrane, whose detailed feedback was instrumental for writing this paper, and to Harald Uhlig for encouraging me to work on a model extension. I also thank Daniel Andrei, Tyler Beason, Hank Bessembinder, Oliver Boguth, Scott Cederburg, Lars Kuehn, Rajnish Mehra, Tobias Sichert and seminar participants at McGill and the Arizona Junior Finance Conference for helpful comments and suggestions.

I. Introduction

In their abstract, Campbell and Cochrane (1999) assert that the habit model "*explains a wide variety of dynamic asset pricing phenomena, including the procyclical variation of stock prices, the long-horizon predictability of excess stock returns, and the countercyclical variation of stock market volatility.*" The fact that these phenomena are captured by a single, intuitive economic mechanism – time-varying risk aversion – has helped to establish the habit model as a central paradigm in asset pricing theory. Campbell and Cochrane's study is covered most asset pricing texts[1] and, at the time of this writing, has accumulated over 6,300 Google Scholar citations.

One limitation of the habit model is the absence of an analytical solution. Wachter (2005) shows that Campbell and Cochrane's numerical solution exhibits inaccuracies that bias the model's implications for *unconditional asset pricing facts*, such as the unconditional mean and volatility of stock market returns. When the model is solved accurately, the original calibration yields a poor fit for these moments. However, Wachter (2005) shows that this issue can be overcome with a different calibration. Surprisingly, neither her study nor subsequent literature has considered the implications of this inaccuracy for *dynamic asset pricing phenomena*, which form the core of Campbell and Cochrane's contribution.

In this paper, I show that the numerical inaccuracy significantly affects dynamics. Most notably, stock market volatility no longer increases during recessions when the model is solved accurately. Instead, it becomes a hump-shaped function of the model's state (the "surplus consumption ratio"). For the worst states, which can be compared to deep recessions in the data, "excess volatility" falls to zero. Instead of being associated with large spikes in stock market volatility, as in the data, recessions therefore feature low volatility in the model.

The counterfactual behavior of volatility turns out to be an inherent property

[1] See, e.g., chapter 8.4 of Campbell et al. (1998), chapter 21 of Cochrane (2005), chapter 10.5.2 of Singleton (2006), chapter 14.2 of Pennacchi (2008), chapter 11.7 of Back (2010), chapter 9.2.3 of Munk (2013), and chapter 6.7.2 of Campbell (2017).

of the habit mechanism, rather than a calibration issue. In particular, time-varying return volatility results exclusively from heteroscedasticity in the price-dividend (P/D) ratio because dividend growth rates are assumed to IID. In turn, the P/D ratio is a function of the surplus consumption ratio, S_t, and therefore inherits its heteroscedasticity from S_t. Campbell and Cochrane specify an exogenous process for the logarithm of S_t whose volatility approaches infinity in the $S_t \to 0$ limit. However, the volatility of S_t's level has to approach zero at that point because log-normal variables cannot turn negative. Economically, a negative surplus consumption ratio would imply that consumption falls below the agent's habit level and render Campbell and Cochrane's utility function undefined. Consequently, the counterfactual behavior of volatility cannot be rectified by recalibrating the model.

I show that the volatility dynamics in the accurate solution have important implications for two other model predictions. First, the model loses its consistency with the "leverage effect", i.e., the empirical observation that returns are negatively correlated with contemporaneous changes in conditional volatility (Black 1976). While the model continues to produce a leverage effect during expansions when solved accurately, the correlation counterfactually becomes positive during recessions. Second, risk premia are less cyclical and returns less predictable than reported by Campbell and Cochrane (1999). This finding reflects two opposing forces. Risk aversion increases during recessions, elevating the volatility of the pricing kernel and, consequently, expected returns. This channel remains unaffected by the numerical solution because the pricing kernel is exogenous. Conversely, return volatility declines during recessions for the accurate solution, lowering expected returns. Overall, the pricing kernel channel dominates, and expected returns remain countercyclical. However, the accurate solution implies that monthly expected returns stay below 15% p.a. during deep recessions, in contrast to the nearly 40% p.a. reported by Campbell and Cochrane (1999). Additionally, in long-horizon predictability regressions based on simulated model data, R-squared values in the original paper are overstated by a factor of about three due to the numerical inaccuracy. At the 7-year horizon, R-squared equals only 10% for the accurate model solution, substantially lower than empirical estimates.

The final section of the paper introduces an extension of the Campbell-Cochrane model based on countercyclical "leverage". I assume, and confirm in the data, that dividend growth becomes more volatile and more correlated with consumption growth in recessions, i.e., for low values of the surplus consumption ratio. Apart from this cyclicality in dividend risk, I calibrate the extended model in exactly the same way as Campbell and Cochrane. The results are striking. The state-dependent characteristics of the price-dividend ratio, expected return, Sharpe ratio, and return volatility in the extended model closely resemble their empirically realistic but inaccurate counterparts in the original paper. Volatility is countercyclical, the leverage effect persists at all times, and excess returns are strongly predictable at long horizons. Hence, the extension "restores" the habit model's realistic predictions for dynamic asset pricing phenomena.

Related Literature. The countercyclical leverage channel resembles the mechanism in Gabaix (2012), who puts forward a model in which dividends' exposure to rare consumption disasters varies over time. In his model, dividends become conditionally more volatile and more correlated with consumption when their disaster exposure rises, and this rise in leverage coincides with falling stock prices. However, leverage does not vary with shocks to the real economy (consumption or dividends) in Gabaix' model and therefore does not rise in recessions. Kuehn et al. (2023) show that this feature leads to counterfactual asset price dynamics during disasters. In reality, leverage increases after a sequence of negative macroeconomic shocks, because the scale of firms' coupon payments to bondholders (financial leverage) as well as the scale of their rent, leasing, and wage payments (operating leverage) cannot easily be adjusted in the short-term. Equity therefore becomes riskier after a sequence of bad shocks. My extension of the habit model captures the responsiveness of leverage to macroeconomic shocks by modelling dividend moments as functions of the surplus consumption ratio.

Bekaert and Engstrom (2017) also extend the habit model, but the nature and objective of their extension are quite different from mine. Their study adds a second state variable that controls persistent variation in the higher moments of endowment shocks. Consumption and dividend growth rates are exposed to identical shocks and therefore perfectly dependent, yet imperfectly correlated due to the non-normal nature of the shocks. These additional non-linearities allows Bekaert and Engstrom to match a broad set of moments of equity index options, in addition to standard equity moments. However, their model is less tractable than Campbell and Cochrane's and features very little return predictability. In contrast, the extension I propose maintains the single state variable, lognormal structure of the original paper and targets the same set of empirical facts, including the long-horizon predictability of excess returns.

Ljungqvist and Uhlig (2015) have previously criticized the habit model on economic grounds, by showing that occasional endowment destructions can be welfare improving for an agent with external habit utility. In response, Campbell and Cochrane (2015) show that this welfare improvement can disappear when the model is solved at daily of higher intervals, depending on how the endowment destruction is spread out over time. The nature of my critique is different. Rather than challenge the economic plausibility of the model's mechanism for rationalizing volatility dynamics and return predictability, my findings cast doubt on the its ability to rationalize these data features. Additionally, I show in the appendix that my results are virtually unchanged when the model is solved at a daily frequency.

II. The Campbell-Cochrane Model

This section briefly describes the model. My notation follows the original paper: Upper case letters are levels, lower case letters are natural logs, and Δ denotes the difference operator; e.g., $\Delta c_{t+1} = \ln\left(\frac{C_{t+1}}{C_t}\right)$.

Consumption follows a homoscedastic random walk,

$$\Delta c_{t+1} = g + v_{t+1} \tag{1}$$

where g is the average log growth rate and $v_{t+1} \overset{\text{IID}}{\sim} N(0, \sigma^2)$. The representative agent's utility function is

$$E_t\left[\sum_{h=0}^{\infty} \delta^h \frac{(C_{t+h} - X_{t+h})^{1-\gamma} - 1}{1-\gamma}\right], \tag{2}$$

where $\delta > 0$ controls time preference and $\gamma > 0$ controls risk preference. Time variation in the habit, X_t, is modelled via the surplus consumption ratio

$$S_t = \frac{C_t - X_t}{C_t}, \tag{3}$$

whose logarithm evolves as a heteroscedastic AR(1) process,

$$s_{t+1} = (1-\phi)\bar{s} + \phi s_t + \lambda(s_t) v_{t+1}. \tag{4}$$

The parameters $\bar{s}$ and ϕ control the mean and persistence of s_t, whereas the function $\lambda(s_t)$ controls its sensitivity to consumption shocks and is specified as

$$\lambda(s_t) = \begin{cases} \bar{S}^{-1}\sqrt{1 - 2(s_t - \bar{s})} - 1 & , s_t < s_{max} \\ 0 & , s_t \geq s_{max} \end{cases}, \tag{5}$$

where $\bar{S} = \sigma\sqrt{\gamma/(1-\phi)}$, $s_{max} = \bar{s} + \frac{1}{2}(1 - \bar{S}^2)$, and $\bar{s} = \ln \bar{S}$. Campbell and Cochrane (1999) show that (5) results in a constant risk-free rate.

The representative agent's intertemporal marginal rate of substitution equals

$$M_{t+1} = \delta e^{-\gamma(\Delta c_{t+1} + \Delta s_{t+1})}. \tag{6}$$

Different from the IMRS in models with recursive utility, (6) does not depend on endogenous objects. The numerical solution method therefore has no effect on the behavior of M_{t+1} in the habit model. This fact will be useful for comparing asset pricing moments across alternative solution methods.

I turn to a discussion of equity valuation. In a Lucas (1978) economy, equity would be modelled as a claim to aggregate consumption. In reality, however, stock market dividends are imperfectly correlated with consumption and considerably more volatile. To capture this "levered" nature of dividends, Campbell and Cochrane (1999) propose to model them as a separate process,

$$\Delta d_{t+1} = g + w_{t+1}, \tag{7}$$

where $w_{t+1} \overset{\text{IID}}{\sim} N(0, \sigma_w^2)$ and $corr(v_t, w_t) = \rho$. Equity represents a claim to the dividends in all future periods. Following Campbell and Cochrane (1999), this approach for modeling leverage has been widely adopted throughout the asset pricing literature.[2] Let $\mathcal{P}(s_t)$ denote the price-dividend ratio, which is a function of the model's state variable s_t. Based on $\mathcal{P}(s_t)$, the cum-dividend return of the dividend claim can be written as

$$R_{t+1} = \frac{\mathcal{P}(s_{t+1}) + 1}{\mathcal{P}(s_t)} e^{\Delta d_{t+1}} \tag{8}$$

and plugged into the standard Euler equation $E_t[M_{t+1} R_{t+1}] = 1$ to yield

$$\mathcal{P}(s_t) = E_t \left[M_{t+1} \left(\mathcal{P}(s_{t+1}) + 1 \right) e^{\Delta d_{t+1}} \right]. \tag{9}$$

The RHS of (9) involves a bivariate integral over v_{t+1} and w_{t+1}. Fortunately, joint normality implies $w_{t+1} | v_{t+1} \sim N\left(\rho \frac{\sigma_w}{\sigma} v_{t+1}, (1-\rho^2)\sigma_w^2 \right)$, which can be used in conjunction with the law of iterated expectations to analytically integrate over w_{t+1}:

$$\mathcal{P}(s_t) = E_t \left[M_{t+1} \left(\mathcal{P}(s_{t+1}) + 1 \right) e^{g + \rho \frac{\sigma_w}{\sigma} v_{t+1} + (1-\rho^2)\sigma_w^2/2} \right] \tag{10}$$

Substituting for M_{t+1} and s_{t+1} from (6) and (4) and writing the expectation as an integral results in

$$\begin{aligned} \mathcal{P}(s_t) =& \delta e^{(1-\gamma)g - \gamma(1-\phi)(\bar{s} - s_t) + (1-\rho^2)\sigma_w^2/2} \\ & \times \int_{-\infty}^{\infty} \left(\mathcal{P}\left((1-\phi)\bar{s} + \phi s_t + \lambda(s_t) v \right) + 1 \right) e^{(\rho \frac{\sigma_w}{\sigma} - \gamma(\lambda(s_t)+1))v} f(v) dv, \end{aligned} \tag{11}$$

where $f(v)$ is the probability density function of a normal random variable with mean zero and standard deviation σ. To examine the model's implications for stock returns in (8), one needs to solve this functional equation for $\mathcal{P}(s_t)$.

[2] Well-known examples include Bansal and Yaron (2004), Routledge and Zin (2010), Gabaix (2012), and Ju and Miao (2012). A predecessor based on Markov-switching fundamentals was proposed by Cecchetti et al. (1993).

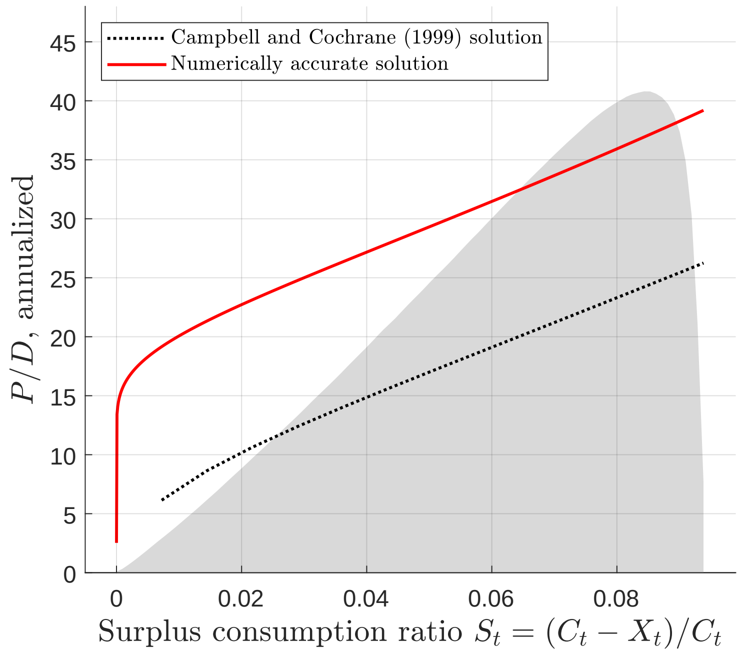

Figure I: Price-dividend ratio. The dotted line replicates Figure 3 in Campbell and Cochrane (1999). The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

Campbell and Cochrane (1999) report results for returns on a claim to consumption in the spirit of Lucas (1978), in addition to those for the claim on dividends. The Euler equation for the consumption claim obtains as a special case of (11) by setting $\rho = 1$ and $\sigma_w = \sigma$. In what follows, I focus on results for the dividend claim because it represents a more realistic counterpart to the stock market.

III. Numerical solution method

Following Campbell and Cochrane (1999), I solve $\mathcal{P}(s_t)$ numerically on a grid by iterating on (11). The only difference between the solution in the original paper and mine lies in the grid values for the state variable s_t, as further discussed below.

On the basis of (11) and a conjectured solution $\mathcal{P}^i(s_t)$, an updated solution

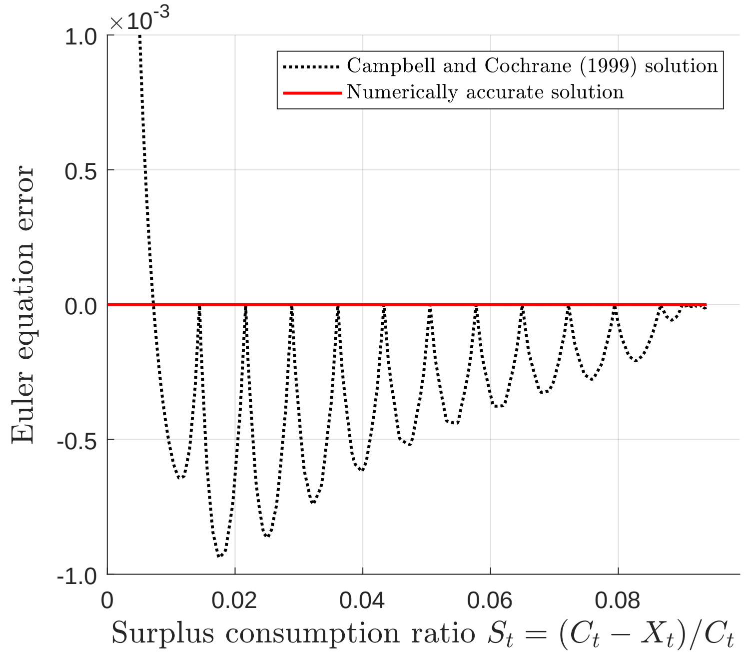

Figure II: Euler equation errors. Euler equation errors are defined in equation (13).

$\mathcal{P}^{i+1}(s_t)$ can be computed as

$$\begin{aligned}\mathcal{P}^{i+1}(s_t) =& \delta e^{(1-\gamma)g-\gamma(1-\phi)(\bar{s}-s_t)+(1-\rho^2)\sigma_w^2/2} \\ &\times \int_{-\infty}^{\infty} \left(\mathcal{P}^i((1-\phi)\bar{s}+\phi s_t+\lambda(s_t)v)+1\right) e^{(\rho\frac{\sigma_w}{\sigma}-\gamma(\lambda(s_t)+1))v} f(v)dv.\end{aligned} \tag{12}$$

This iterative procedure is repeated until the supremum norm of $\mathcal{P}^{i+1} - \mathcal{P}^i$ falls below 10^{-8}. I evaluate the integral numerically using Gauss-Chebyshev quadrature with 50 nodes, spread between -7 and $+7$ standard deviations. To compute $\mathcal{P}$ for off-grid values of s_{t+1} that result from the quadrature, I linearly interpolate $\ln \mathcal{P}$ as a function of s_{t+1}. Parameter values for the model's monthly calibration are identical to those in Campbell and Cochrane (1999) and shown in Table 1 of their paper.

Campbell and Cochrane do not specify the construction of their state grid, but the values can be gleaned from Figure A1 of their online appendix and are also reported by Wachter (2005). The grid is based on 14 equally-spaced points for S_t between 0 and $e^{s_{max}}$, excluding 0 (netting 13 points), as well as points at 0.090, 0.091, 0.092, and 0.093, for a total of 17 points. The dotted line in Figure I shows the numerical solution for the price-dividend ratio $\mathcal{P}(s_t)$ based on this sparse grid, which closely matches its counterpart in Figure 3 in Campbell and Cochrane (1999). Similarly, the first two conditional returns moments in Figures III and V below match their counterparts in Figures 5 and 4 in Campbell and Cochrane (1999). The numerical solution based on the sparse state grid therefore successfully replicates the results in the original paper.

Figure II assesses the accuracy of Campbell and Cochrane's numerical solution by plotting Euler equation errors, defined as

$$\varepsilon(s_t) = E_t \left[M_{t+1} \frac{\mathcal{P}(s_{t+1}) + 1}{\mathcal{P}(s_t)} e^{\Delta d_{t+1}} \right] - 1. \tag{13}$$

I compute (13) for both on-grid and off-grid values of s_t and rely on the same quadrature and interpolation method as in the model solution. By construction, Euler equation errors are close to zero for grid values of S_t. For off-grid values, however, they are large. Absolute errors have an average of 2.7×10^{-4} and a maximum of 0.21, which indicates that the numerical solution is fairly imprecise.

Wachter (2005) shows that an accurate solution of the habit model crucially relies on a grid with more S_t values close to zero. For expositional simplicity, I use a different algorithm than Wachter (2005), but it is important to emphasize that my contribution relative to her study lies in an examination of different model implications, not in proposing a superior algorithm.[3] I use a state grid with $n_s = 100,000$ points for s_t that contains more gridpoints for small values of s_t and fewer points for large values. Specifically, the i-th gridpoint equals $s_{min} + \mathcal{D} \times (i-1)^{\pi}$, where $s_{min} = -300$, $\mathcal{D} = (s_{max} - s_{min})/(n_s - 1)^{\pi}$, and $\pi = 1/10$. Figure II shows that the solution based on this state grid features much smaller Euler equation errors than the original solution. Absolute Euler equation errors have an average of 2.4×10^{-9} and a maximum of 5.1×10^{-8}. Hence, the solution is very precise. Obviously, the computational costs of using such a fine grid would have been prohibitive when

[3] I find that the preferred solution approach of Wachter (2005), "Grid 3, Series method", results in absolute Euler equation errors that have an average of 3.9×10^{-6} and a maximum of 5.2×10^{-4}. This is slightly less accurate than my solution, but considered accurate by common standards.

Campbell and Cochrane wrote their paper in the mid 1990s.

The solid line in Figure I shows that the resulting P/D ratio looks drastically different from the original solution. The reason for this difference is also immediately clear: P/D features extreme curvature for small values of S_t and approaches zero for $S_t \to 0$. This "boundary condition" is a critical feature of the habit mechanism: As consumption approaches the agent's habit level (as S_t approaches zero), risk aversion approaches infinity, expected returns approach infinity, and the price-dividend ratio therefore approaches zero. Campbell and Cochrane's solution fails to capture the curvature in P/D because their lowest grid value of $S_t = 0.0072$ lies above the concave region of P/D. Despite the fact that states to the left of this point occur with a probability of just 0.5%, they have a sizable effect on the level of P/D in regions with more probability mass due to the recursive definition of P/D.

The appendix contains several robustness tests. First, I solve $\mathcal{P}(s)$ based on a projection method (Judd 1992), as an alternative to the finite element method described above. Global polynomial approximations of this kind have previously been shown to accurately capture nonlinearities in asset pricing models with long-run risks (Pohl et al. 2018). Second, I solve $\mathcal{P}(s)$ based on Wachter's (2005) "series method", i.e., by computing a series of dividend strip prices rather than solving for the fixed point of a functional equation. Third, I solve $\mathcal{P}(s)$ based on a daily rather than monthly decision interval. This approach follows Campbell and Cochrane (2015), who use daily simulations to approximate the continuous time limit of the habit model. In each case, the appendix shows that both $\mathcal{P}(s)$ and the conditional moments of returns are almost indistinguishable from their counterparts in the benchmark solution.

A careful reader may notice that $\mathcal{P}(s)$ in Figure I does not completely reach its theoretical limit of zero as $S_t \to 0$ (equivalently, as $s_t \to -\infty$).[4] This effect is caused by $\mathcal{P}(s)$ approaching its limit very slowly as $s_t \to -\infty$, and the lowest grid

[4]Similarly, the conditional volatility of returns in Figure III does not completely reach its theoretical limit of σ_d as $S_t \to 0$.

value of -300 for s_t not being "close enough" to $-\infty$. In a forth robustness test, I show that a projection approach based on a very high dimensional polynomial is able to capture $\mathcal{P}(s)$ far enough into the left tail to make it visually indistinguishable from zero. I do not use this case as my benchmark because (i) it is extremely inefficient computationally, (ii) it makes no discernible difference for asset pricing moments in the region $s_t > -300$, and (iii) it requires a daily decision interval, which makes it harder to illustrate why my results differ from those in the original study. Importantly, however, this robustness test illustrates that, in order to find accurate return moments in the economically relevant region of the state space, it is not necessary to compute $\mathcal{P}(s)$ all the way to the left limit of its domain.

IV. Volatility Dynamics

Figure III shows the conditional volatility of returns as a function of the surplus consumption ratio. Because returns are a function of the P/D ratio, their volatility is affected by the numerical solution. The dotted line is based on the sparse state grid and replicates Figure 5 in Campbell and Cochrane (1999). It shows that, as the surplus consumption ratio falls, the conditional volatility of returns rises monotonically. Since low values of S_t are also associated with low P/D ratios, Campbell and Cochrane (1999) naturally interpret them as recessions and highlight the "countercyclical" nature of volatility as one of the model's main predictions.

The solid line shows that volatility is a hump-shaped (rather than monotonic) function of S_t for the numerically accurate model solution. For values below the 18.2-th percentile of S_t, which is indicated by the vertical line in Figure III, $\sigma_t[R_{t+1}]$ is an increasing function of S_t. The probability of these states roughly lines up with the NBER recession frequency of 13.6% in the 1947-1922 sample. Hence, the model predicts that recessions are associated with below peak volatility and that deep recessions, such as the financial crisis of 2008, are associated with particularly low volatility. This prediction runs opposite to the data, where recessions are typically associated with large spikes in stock market volatility.

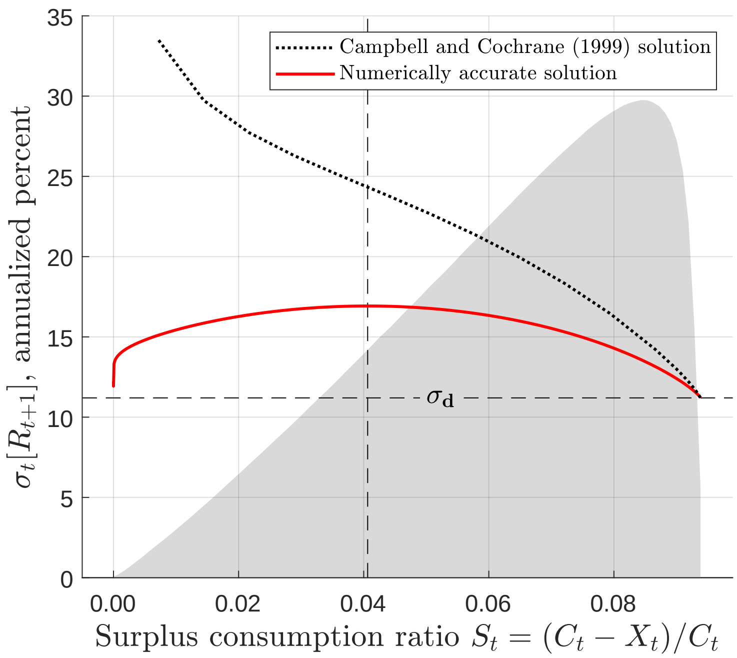

Figure III: Conditional volatility of returns. The dotted line replicates Figure 5 in Campbell and Cochrane (1999). The dashed vertical line marks the point at which volatility reaches its maximum for the numerically accurate model solution. The dashed horizontal line marks the volatility of log dividend growth rates. The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

The results in Figure III also affect the habit model's ability to capture the "leverage effect", i.e., the empirical observation that returns are negatively correlated with contemporaneous changes in conditional volatility (Black 1976). Campbell and Cochrane (1999) (p222) state that their model captures the leverage effect since equity valuations (the P/D ratio) are an increasing function of S_t, whereas volatility is a *decreasing* function of S_t in their solution. For the numerically accurate solution, however, volatility is an *increasing* function of S_t in recessions, which implies that the correlation between returns and volatility innovations becomes positive. In contrast, the data shows that the leverage effect remains significantly negative during recessions. In the 1947-2022 sample, the correlation equals -45.1% during NBER recession months and -43.6% outside of NBER recession months.[5]

[5] I compute correlations based on an overlapping sample of 21 trading day returns on the CRSP value-weighted market index. The conditional volatility of the return between day t and day $t+21$

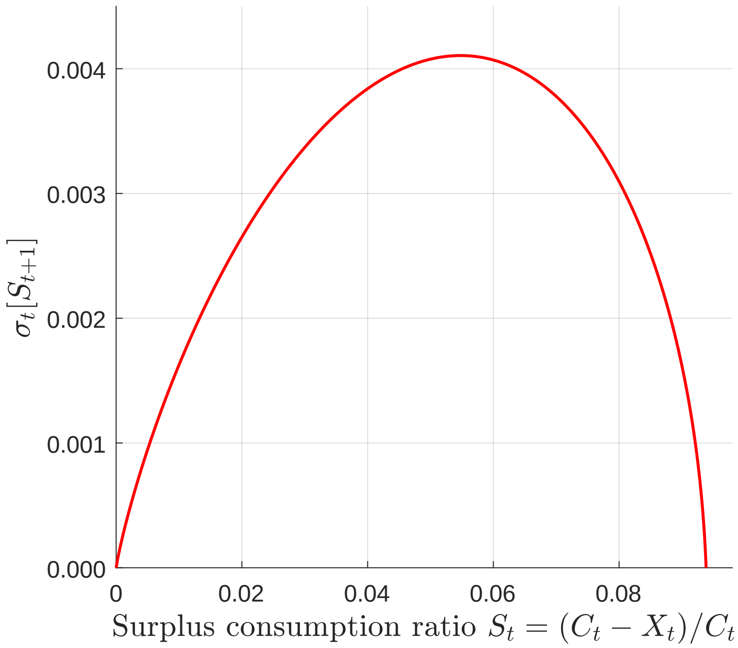

Figure IV: Conditional volatility of the surplus consumption ratio.

What accounts for the hump-shape of volatility? Because dividends are homoscedastic, time variation in the volatility of returns is entirely driven by the heteroscedasticity of P/D, which in turn reflects heteroscedasticity in the surplus consumption ratio. Standard results for log-normal random variables imply

$$\sigma_t(S_{t+1}) = E_t[S_{t+1}]\sqrt{e^{\mathrm{Var}_t[s_{t+1}]} - 1}, \tag{14}$$

where $E_t[S_{t+1}] = e^{(1-\phi)\bar{s}+\phi s_t+\lambda(s_t)^2\sigma^2/2}$ and $e^{\mathrm{Var}_t[s_{t+1}]} = e^{\lambda(s_t)^2\sigma^2}$. Figure IV shows that $\sigma_t(S_{t+1})$ is a hump-shaped function of S_t and equal to zero at the endpoints of the domain $[0, e^{s_{max}}]$. Because S_t has a conditional volatility of zero at the endpoints of its domain, so does P/D. As a result, the conditional volatility of returns approaches the volatility of dividends at these points, which is time-invariant and indicated by the horizontal line in Figure III.

is estimated by the standard deviation of daily returns between days $t-22$ and $t-1$. To assess statistical significance, I regress changes in the 21-day volatility on the contemporaneous return. Using Newey-West standard errors with 21 lags to account for the overlapping data, I find t-statistics of -3.4 and -7.8 for the recessions and non-recession subsample, respectively. The leverage effect is therefore significantly negative across the business cycle.

Importantly, the fact that $\sigma_t(S_{t+1})$ approaches zero at both 0 and $e^{s_{max}}$ is necessary in order for habit utility to be mathematically defined. For $S_t > e^{s_{max}}$, the sensitivity function $\lambda(s_t)$ in (5) turns negative, which implies that the log surplus consumption ratio s_t in (4) has a negative conditional volatility. For $S_t < 0$, consumption falls below the habit level and utility in (2) is not well defined. To prevent S_t from falling outside of the interval $[0, e^{s_{max}}]$ with a positive probability, the volatility of S_t has to approach zero at its endpoints. Hence, this feature is a mathematical necessity rather than merely a property of a particular calibration.

V. Return Predictability

The dynamics of volatility are important for the dynamics of expected returns. Specifically, because the accurate solution produces lower (systematic) return volatility in recessions, it also produces a lower equity premium in these states. Figure V shows expected returns as a function of the surplus consumption ratio. The dotted line is based on the sparse state grid and replicates Figure 4 in Campbell and Cochrane (1999). It shows that expected returns can spike to nearly 40% p.a. during recessions, i.e., for low values of the surplus consumption ratio. In contrast, the solid line for the numerically accurate solution shows that expected returns remain below 15% p.a. for the same states.[6] Risk premia in the habit model are therefore considerably less countercyclical than suggested by Campbell and Cochrane (1999).

To understand time-variation in expected returns in more detail, it is useful to write them as[7]

$$E_t[R_{t+1}] = R^f - \frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}\sigma_t[R_{t+1}]\rho_t[M_{t+1}, R_{t+1}]. \tag{15}$$

I discuss each RHS term in equation (15). First, R^f in time-invariant by construction and therefore doesn't affect the cyclicality of expected returns. Second,

[6]Of course, risk aversion approaches infinity in the $S_t \to 0$ limit and expected returns therefore do so as well, but these states occur with a negligible probability. Based on a long simulation, the probability of expected returns exceeding 15% p.a. equals 0.0002%.

[7]$E_t[M_{t+1}R_{t+1}] = 1$ implies that the equity premium equals $E_t[R_{t+1}] - R^f = -\frac{cov_t[M_{t+1}, R_{t+1}]}{E_t[M_{t+1}]}$. Replacing the covariance by $\sigma_t[M_{t+1}]\sigma_t[R_{t+1}]\rho_t[M_{t+1}, R_{t+1}]$ and re-arranging yields (15).

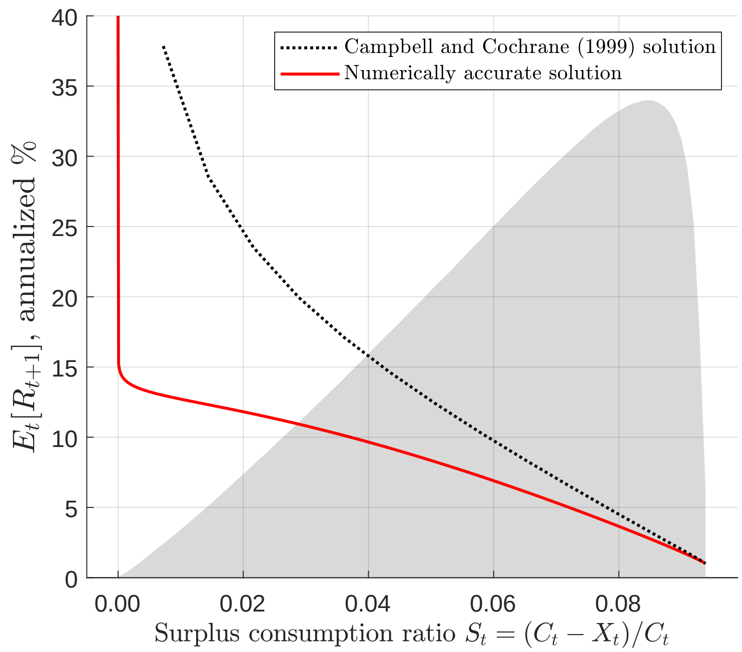

Figure V: Expected returns. The dotted line replicates Figure 4 in Campbell and Cochrane (1999). The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

$\frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}$ equals the maximum Sharpe ratio, as implied by the Hansen and Jagannathan (1991) bound. Because the pricing kernel is exogenous, the maximum Sharpe ratio is identical for both model solutions. Figure VI shows that $\frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}$ is monotonically decreasing in S_t and approaches infinity in the $S_t \to 0$ limit, where risk aversion becomes unbounded. Unsurprisingly, time-varying risk aversion therefore contributes towards increasing the equity premium in recessions, and this effect is not altered by the accurate solution. Third, in contrast to the results reported by Campbell and Cochrane (1999), $\sigma_t[R_{t+1}]$ is an increasing function of S_t in recessions for the accurate solution. Equation (15) shows that return volatility therefore contributes towards decreasing the equity premium during recessions, which makes risk premia less cyclical for the accurate solution. Fourth, to see how the correlation between returns and the pricing kernel affects the cyclicality of expected returns, note that equation (15) implies $\rho_t[M_{t+1}, R_{t+1}] = -\left(\frac{E_t[R_{t+1}]-R^f}{\sigma_t[R_{t+1}]}\right) / \left(\frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}\right)$. This is a

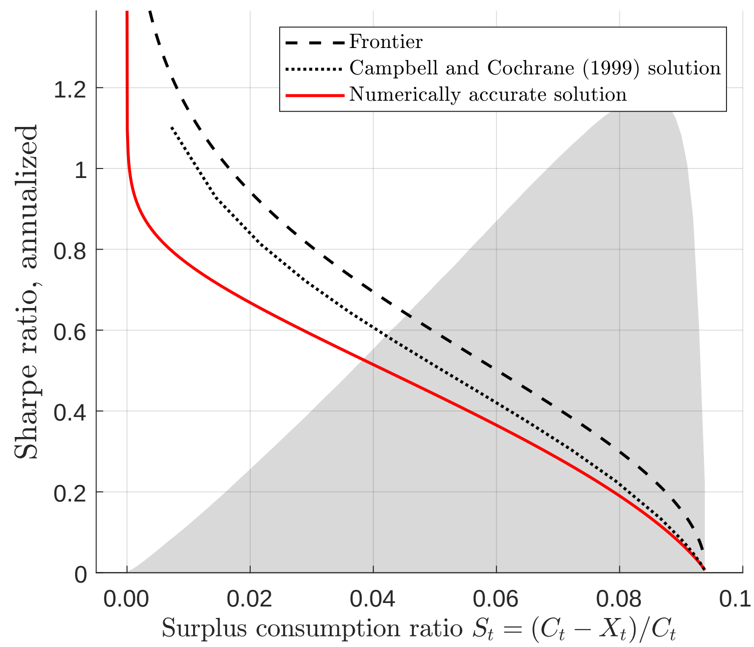

Figure VI: Sharpe ratios. The maximum Sharpe ratio $\frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}$ (dashed line) is identical for both model solutions because the pricing kernel is exogenous. The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

familiar result: If the correlation is closer to -1, the stock market's Sharpe ratio is closer to Hansen Jagannathan bound. Figure VI shows that the difference between the maximum Sharpe ratio and that of the stock market is larger for the accurate solution than that of Campbell and Cochrane (1999), and particularly so during recessions. Hence, the correlation between returns and the pricing kernel is less negative during recessions for the accurate solution, and this channel also contributes towards making expected returns less cyclical than reported in the original paper.

The results up to this point have focused on one-period returns. However, Campbell and Cochrane (1999) primarily emphasize the habit model's ability to produce variation in long-horizon returns, because this variation can be quantified empirically based on predictive regressions. Of course, the dynamics of long-horizon expected returns are closely related the dynamics of one-period expected returns. Figure VII

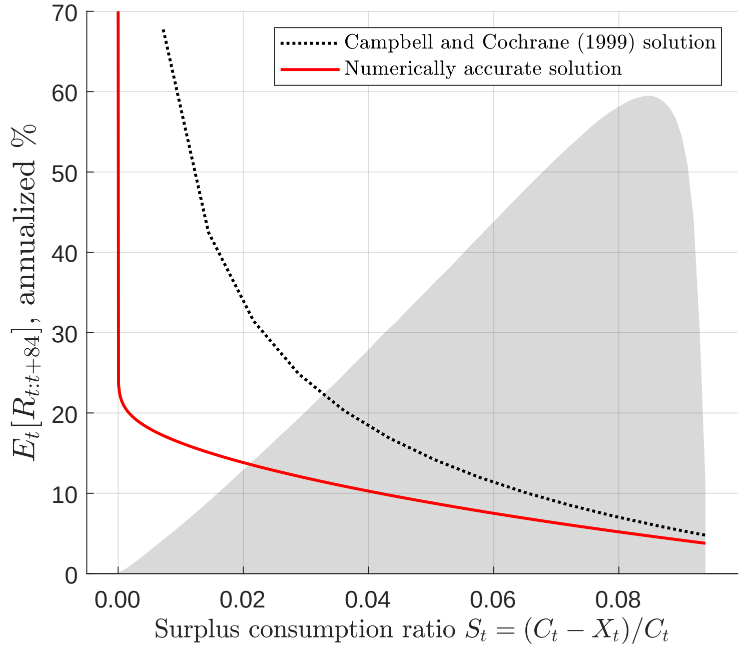

Figure VII: Long-horizon expected returns. 7-year expected returns as a function of the surplus consumption ratio. The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

shows 7-year expected returns as a function of the surplus consumption ratio.[8,9] Similar to one-period expected returns in Figure V, long-horizon expected returns rise considerably less during recessions for the accurate solution. They can spike to over 65% p.a. during recessions for the Campbell and Cochrane (1999) solution, but stay below 20% p.a. in the same states for the accurate solution.

Because long-horizon returns are less cyclical for the accurate solution, they are also less predictable by cyclical variables. Table I shows estimates from regressions of log excess returns on the log P/D ratio for various horizons. To compute these statistics, I follow Campbell and Cochrane (1999) in simulating a long history of artificial data from the model and time-aggregating them to an annual frequency.

[8]Seven years corresponds to the longest horizon Campbell and Cochrane (1999) consider in their predictability regressions.

[9]I compute long-horizon expected returns recursively as $E_t[R_{t:t+h}] = E_t\left[\frac{\mathcal{P}(s_{t+1})+1}{\mathcal{P}(s_t)} e^{\Delta d_{t+1}} E_{t+1}[R_{t+1:t+h}]\right]$ and rely on the same interpolation method and quadrature routine as in the model solution to evaluate the outer expectation.

Table I: Long horizon return regressions

The two data columns are taken from Table 5 in Campbell and Cochrane (1999), "Campbell-Cochrane (1999) solution" replicates the inaccurate model results in the original paper, and "Numerically accurate solution" is based on the fine grid described in Section III.

	Postwar sample		Long sample		Campbell-Cochrane (1999) solution		Numerically accurate solution	
Horizon (Years)	10× Coefficient	R^2	10× Coefficient	R^2	10× Coefficient	R^2	10× Coefficient	R^2
1	-2.6	0.18	-1.3	0.04	-1.8	0.07	-1.4	0.03
2	-4.3	0.27	-2.8	0.08	-3.4	0.13	-2.6	0.05
3	-5.4	0.37	-3.5	0.09	-4.9	0.18	-3.7	0.06
5	-9.0	0.55	-6.0	0.18	-7.1	0.26	-5.5	0.09
7	-12.1	0.65	-7.5	0.23	-8.9	0.30	-6.8	0.10

The two data columns are taken from Table 5 in Campbell and Cochrane (1999), and the column "Campbell-Cochrane (1999) solution" replicates their (inaccurate) results for the model: The magnitudes of both R^2 values and regression slope coefficients are steeply increasing in the return horizon. At the 7-year horizon, the R^2 reaches 30% and is quantitatively comparable to the empirical estimates. For the accurate solution, slope coefficients and R^2 values are both smaller in magnitude. Relative to the inaccurate solution, the R^2 for the 7-year horizon falls by a factor of 3 to a value of 10%, which is substantially smaller than the empirical estimates.

In sum, when the habit model is solved accurately, it implies falling volatility during recessions, which makes it inconsistent with the leverage effect and the strong predictability of excess returns we observe in the data.

VI. A Habit Model with Countercyclical Leverage

I now propose an extension of the Campbell and Cochrane (1999) model, which captures the dynamic asset pricing phenomena the original paper sought to explain. The extension intentionally remains as close as possible to the original model specification and targets the same empirical phenomena, rather than adding complexity and targeting a wider set of moments. Specifically, the extended model features the surplus consumption ratio as its sole state variable and assumes lognormal shocks, which implies that many findings about the habit model in existing work will likely

Table II: Countercyclical leverage in the data

Δc is the log growth rate of real annual nondurables and services consumption per capita. Δd is the annual log growth rate of CRSP cash dividends, converted to real units based on the implicit consumption deflator. Calendar years with at least six NBER recession months are classified as recessions. The post-war sample spans 1947-2022 and the long sample spans 1929-2022.

	Post-war sample			Long sample		
	All years	Recession	Non-recession	All years	Recession	Non-recession
$\sigma[\Delta d]/\sigma[\Delta c]$	4.36	7.16	4.10	4.71	5.12	4.45
$\rho[\Delta d, \Delta c]$	0.12	0.31	-0.01	0.49	0.64	0.31

continue to hold. Novel assumptions and associated empirical evidence is discussed in Section A, the calibration in Section B, and asset pricing results in Section C.

A. Assumptions

The volatility of dividends and the correlation between dividends and consumption are modelled as functions of the surplus consumption ratio,

$$\begin{aligned} \sigma_d(s_t) =& h_0 + h_1\sqrt{1-2(s_t-\bar{s})} \\ \rho(s_t) =& min\left\{1, r_0 + r_1\sqrt{1-2(s_t-\bar{s})}\right\}, \end{aligned} \tag{16}$$

where the parameters (h_0, h_1, r_0, r_1) replace (σ_d, ρ) in the original specification. All other model elements remain unchanged. For $h_1 > 0$, dividend volatility is decreasing in s_t, so that recessions (states of low s_t) are associated with elevated volatility. For $r_1 > 0$, consumption and dividends become more correlated in bad times.

To evaluate whether the assumed time-variation is empirically plausible, I compute it in annual consumption and dividend data. Consumption equals real nondurables and services consumption per capita from the BEA. Cash dividends on the value-weighted stock market index are taken from CRSP and converted to real units based on the implicit consumption deflator. Calendar years are classified as recession years if they contain at least six NBER recession months and as non-recession years otherwise. I then compute the ratio of dividend growth volatility to consumption growth volatility and the correlation between consumption and dividend growth separately for recession and non-recession years. Table II shows the results. In the 1947-2022 postwar sample, the dividend-to-consumption volatility ratio rises from

4.10 in expansions to 7.16 in recessions, whereas the correlation between consumption and dividends rises from -0.01 in expansions to 0.31 in recessions. Results for the longer 1929-2022 sample look similar. The dividend-to-consumption volatility ratio rises from 4.45 in expansions to 5.12 in recessions, whereas the correlation between consumption and dividends rises from 0.31 in expansions to 0.64 in recessions. For $h_1 > 0$ and $r_1 > 0$, the dynamics in equation (16) are therefore qualitatively consistent with the countercyclical nature of dividend risks in the data.[10]

B. Calibration

I solve the model based on the accurate solution method described above and rely on the original calibration, apart from the parameters that govern $\sigma_d(s_t)$ and $\rho(s_t)$. The solution is once again very precise: Absolute Euler equation errors have an average of 1.9×10^{-9} and a maximum of 6.8×10^{-9}.

The leverage parameters ($h_0 = 0.086, h_1 = 0.028, r_1 = 0.011, r_1 = 0.2$) are calibrated to an unconditional volatility of 11.2% p.a. for dividends and an unconditional correlation of 0.2 between consumption and dividends, as in Campbell and Cochrane's original calibration. The remaining two degrees of freedom are used to generate realistic asset price dynamics, as illustrated below. Figure VIII shows $\sigma_d(s_t)/\sigma$ and $\rho(s_t)$ as functions of the surplus consumption ratio. To compare the amount of time variation in these moments to the data, I compute them in simulated model data that are time-aggregated to an annual frequency. I classify year t as a recession year if S_t falls below its 13.6-th percentile (the frequency of NBER recessions in postwar data) at the end of year $(t-1)$. My calibration implies that the ratio of dividend volatility to consumption volatility rises from 7.3 in expansion years to 9.0 in recession years, whereas the correlation between consumption and dividends rises from 0.17 in expansions to 0.33 in recessions. The amount of time variation in both of these moments is comparable its data analogue in Table II. At

[10] Schreindorfer (2020) documents the related fact that dividend growth rates (and returns) are more correlated with consumption growth rates conditional on a low consumption realization, i.e., their "downside correlation" increases in the left tail.

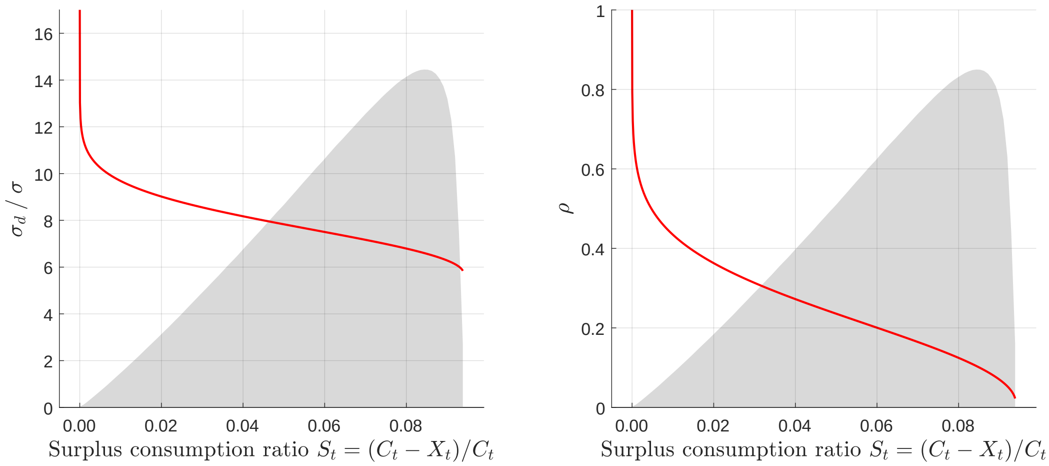

Figure VIII: Extended habit model: Countercyclical leverage. The ratio of dividend volatility to consumption volatility (left panel) and the correlation between consumption and dividend growth (right panel) for the extended habit model are plotted as functions of the surplus consumption ratio. The unconditional distribution of the surplus consumption ratio (not drawn to scale) is shown in the background.

the same time, the level of both moments differs somewhat from the data because, as mentioned above, it is calibrated to the values in Campbell and Cochrane (1999).

C. Results

The extended model mimics the empirically realistic but inaccurate asset price dynamics reported by Campbell and Cochrane (1999). Specifically, Figure IX shows that the price dividend ratio (top-left panel), return volatility (top-right panel), Sharpe ratio (bottom-left panel), and 7-year expected returns (bottom-right panel) in the accurately-solved extended model (solid line) closely match the corresponding moments in the inaccurately-solved original model (dotted line): Volatility is a monotonically decreasing function of S_t and therefore increasing in recessions. As a result, returns are negatively correlated with volatility innovations at all times, which implies that the model captures the leverage effect. Furthermore, expected returns are steeply decreasing in S_t. In long-horizon predictability regressions, the slope coefficient rises from -1.8$\times 10^{-1}$ at the 1-year horizon to -8.4$\times 10^{-1}$ at the 7-year horizon, while R-squared rises from 0.06 to 0.24 (not tabulated). These values

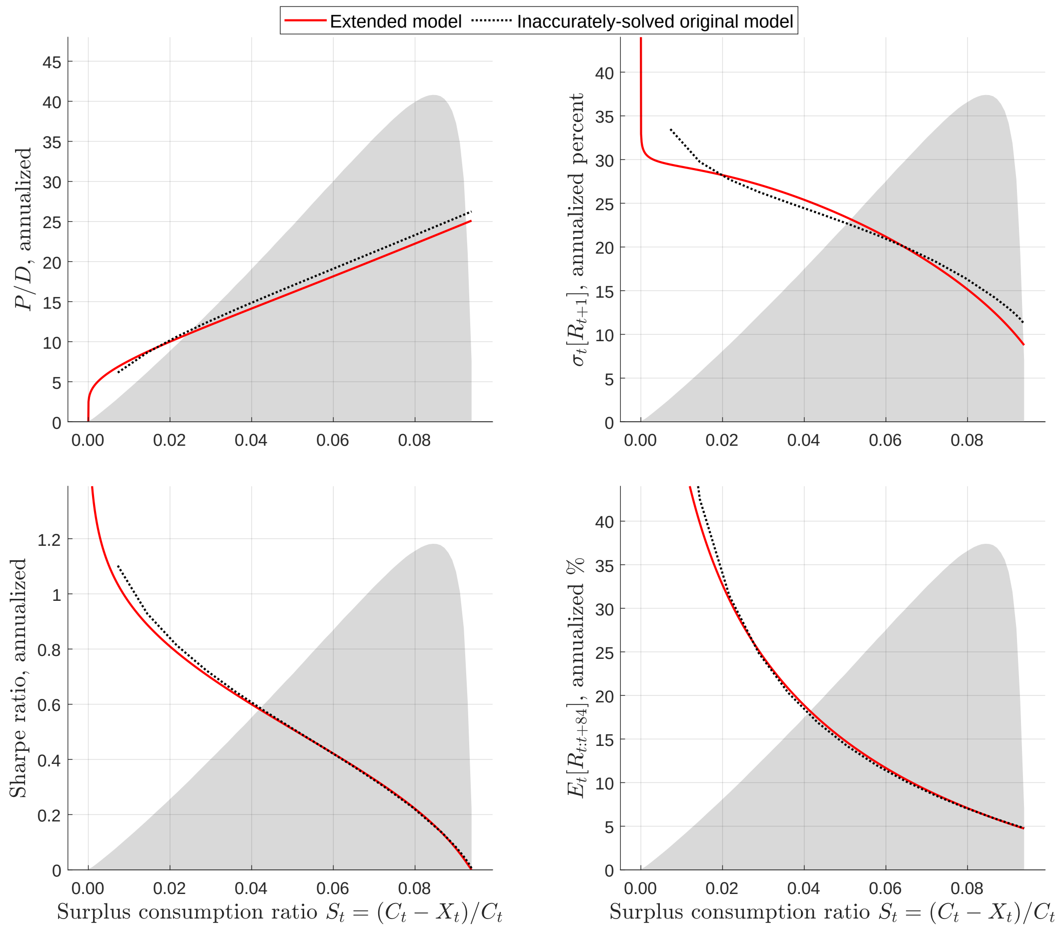

Figure IX: Extended habit model: Asset prices. This figure shows asset pricing moments for the extended habit model and compares them to the equivalent moments in the inaccurately-solved version of Campbell and Cochrane (1999).

are close to those for the inaccurate solution of the original habit model in Table I and therefore close to the data.

Why does the countercyclical leverage model succeed where the accurately-solved original model did not? First, as discussed in Section IV, the original model implies that return volatility falls during recessions due to a decline in the conditional volatility of P/D. In the countercyclical leverage model, P/D volatility continues to decline during recessions and it continues to approach zero for $S_t \to 0$ (not shown). However, dividends become more volatile at the same time and this effect dominates the reduction in P/D volatility, making returns more volatile overall. Second, as discussed in Section V, recessions do not feature large spikes in expected returns in the original model because (i) return volatility falls and (ii) the correlation between returns and the pricing kernel becomes less negative. In the extended model, return volatility does not decrease in recessions, and dividends and consumption become more correlated. Returns (which reflect dividends) and the pricing kernel (which reflects consumption) therefore become more (negatively) correlated as well, which increases the Sharpe ratio and results in larger spikes in expected returns.

VII. Conclusion

When solved accurately, the Campbell and Cochrane (1999) model implies that volatility is a hump-shaped function of the surplus consumption ratio. Deep recessions are therefore characterized by very low volatility, as opposed to large volatility spikes, as in the data. These volatility dynamics are an inherent feature of the habit mechanism, rather than an implication of a specific calibration. In the accurate solution, the "leverage effect" counterfactually disappears in recessions and stock market returns are substantially less predictable than suggested by the results in Campbell and Cochrane (1999).

I show that these undesirable implications can be overturned by augmenting the habit model with countercyclical leverage – a robust feature of dividend growth rates in the data. Intuitively, when leverage rises in recessions, it makes stocks riskier and therefore increases their expected return. Most representative agent models follow Cecchetti et al. (1993), Campbell and Cochrane (1999), and Abel (1999) in modelling leverage as time-invariant, and therefore abstract from an important driver of time-varying stock market risk in the data. It appears likely that countercyclical leverage would result in more realistic asset price dynamics in other model frameworks as well, and it would be interesting to explore this possibility in future work.

References

Abel, A. B. (1999): "Risk premia and term premia in general equilibrium," *Journal of Monetary Economics*, 43, 3–33.

Back, K. (2010): *Asset Pricing and Portfolio Choice Theory*, Oxford University Press.

Bansal, R. and A. Yaron (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *Journal of Finance*, 59, 1481–1509.

Bekaert, G. and E. Engstrom (2017): "Asset return dynamics under habits and bad environment–good environment fundamentals," *Journal of Political Economy*, 125, 713–760.

Black, F. (1976): "Studies of stock price volatility changes," *Proceedings of the 1976 Meetings of the American Statistical Association*, 171–181.

Campbell, J. (2017): *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton University Press.

Campbell, J. Y. and J. H. Cochrane (1999): "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, 107, 205–251.

——— (2015): "The fragile benefits of endowment destruction," *Journal of Political Economy*, 123, 1214–1226.

Campbell, J. Y., A. W. Lo, A. C. MacKinlay, and R. F. Whitelaw (1998): "The econometrics of financial markets," *Macroeconomic Dynamics*, 2, 559–562.

Cecchetti, S. G., P.-s. Lam, and N. C. Mark (1993): "The equity premium and the risk-free rate: Matching the moments," *Journal of Monetary Economics*, 31, 21–45.

Cochrane, J. (2005): *Asset Pricing: Revised Edition*, Princeton University Press.

Gabaix, X. (2012): "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance," *The Quarterly journal of economics*, 127, 645–700.

Hansen, L. P. and R. Jagannathan (1991): "Implications of security market data for models of dynamic economies," *Journal of Political Economy*, 99, 225–262.

Ju, N. and J. Miao (2012): "Ambiguity, learning, and asset returns," *Econometrica*, 80, 559–591.

Judd, K. L. (1992): "Projection methods for solving aggregate growth models," *Journal of Economic theory*, 58, 410–452.

Kuehn, L.-A., D. Schreindorfer, and F. Schulz (2023): "Persistent Crises and Levered Asset Prices," *The Review of Financial Studies*, 36, 2571–2616.

Ljungqvist, L. and H. Uhlig (2015): "Comment on the Campbell-Cochrane Habit Model," *Journal of Political Economy*, 123, 1201–1213.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Munk, C. (2013): *Financial Asset Pricing Theory*, Oxford University Press.

Pennacchi, G. G. (2008): *Theory of asset pricing*, Pearson/Addison-Wesley Boston.

Pohl, W., K. Schmedders, and O. Wilms (2018): "Higher order effects in asset pricing models with long-run risks," *The Journal of Finance*, 73, 1061–1111.

Routledge, B. R. and S. E. Zin (2010): "Generalized disappointment aversion and asset prices," *The Journal of Finance*, 65, 1303–1332.

Schreindorfer, D. (2020): "Macroeconomic tail risks and asset prices," *The Review of Financial Studies*, 33, 3541–3582.

Singleton, K. (2006): *Empirical Dynamic Asset Pricing: Model Specification and Econometric Assessment*, Princeton University Press.

Wachter, J. A. (2005): "Solving models with external habit," *Finance Research Letters*.