North American Journal of Economics and Finance 40 (2017) 55–62

Contents lists available at ScienceDirect

North American Journal of Economics and Finance

journal homepage: www.elsevier.com/locate/ecofin

Campbell and Cochrane meet Melino and Yang: Reverse engineering the surplus ratio in a Mehra–Prescott economy ☆

Jim Dolmas *

Federal Reserve Bank of Dallas, 2200 North Pearl Street, Dallas, TX 75201, United States

A R T I C L E I N F O

Article history:
Received 20 August 2015
Received in revised form 19 January 2017
Accepted 23 January 2017

JEL classification:
E44
G12

Keywords:
Habit
Asset returns
Stochastic discount factor
State-dependent preferences

A B S T R A C T

The well-known habit model of Campbell and Cochrane (1999) specifies a process for the 'surplus ratio'—the excess of consumption over habit, relative to consumption—rather than an evolution for the habit stock. This paper shows that Campbell–Cochrane preferences can be accommodated in a Markov chain framework à la Mehra and Prescott (1985) and mapped into an equivalent state-dependent form of the sort studied by Melino and Yang (2003). The equivalence sheds light on the workings of Campbell–Cochrane preferences and the plausibility of upcounting in Melino and Yang's framework. The result may also have some pedagogical value.

1. Introduction

In this note, I demonstrate one way of putting the habit preferences of Campbell and Cochrane (1999) into the two-state Markov chain framework of Mehra and Prescott (1985). I expect a natural question in the minds of at least a few readers is, "Why?" The answer is that by situating Campbell–Cochrane preferences in a Mehra–Prescott economy we can perform another sort of 'reverse engineering' exercise, complementary to that performed by Campbell and Cochrane themselves. The reverse engineering draws on the work of Melino and Yang (2003), who showed us, in the context of a Mehra–Prescott economy, exactly what the stochastic discount factor (SDF) must look like to match the first and second moments of asset returns in Mehra and Prescott's long sample of returns. We can calibrate our version of Campbell–Cochrane preferences to match that SDF.

The exercise would be of only pedagogical interest unless it told us something interesting about one or both of the two approaches to the equity premium puzzle that it combines. I think it does. While countercyclical risk aversion has been rightly emphasized as a key mechanism in the Campbell–Cochrane model, mapping Campbell and Cochrane into a state-dependent preference specification that matches the returns data shows that a countercyclical utility discount factor, often

☆ The views expressed herein are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

* Corresponding author.
E-mail address: jim@jimdolmas.net

http://dx.doi.org/10.1016/j.najef.2017.01.006

greater than one, is also important. And, while Melino and Yang dismiss state-dependent specifications that imply discount factors greater than one, the model here shows there may be a plausible story that rationalizes such a specification.

That said, the exercise may have pedagogical value as well. Given the computational tractability of the Markov chain framework, and the ubiquity of analogous structures in macroeconomics, the framework is a natural one in which to teach asset-pricing within a graduate course in macroeconomics.[1] Most of the major responses to the 'equity premium puzzle' fit easily within the framework—except for Campbell and Cochrane.[2] The exercise in this paper fills that gap.

It is useful to quickly review the features of the Campbell–Cochrane and Melino–Yang models separately before combining them. The next two sections do this.

1.1. Campbell and Cochrane

Campbell and Cochrane's 1999 paper in the *Journal of Political Economy* employs habit formation to successfully resolve a number of asset pricing puzzles, including Mehra and Prescott's equity premium puzzle. Campbell and Cochrane achieve these resolutions by a clever reverse engineering of their representative agent's habit process.

Rather than specify a law of motion for the habit stock, Campbell and Cochrane specify a law of motion for what they call the 'surplus ratio', $S_t = (c_t - h_t)/c_t$, where c_t is aggregate consumption (the habit is external) and h_t is the habit stock. Their stochastic discount factor, from t to $t+1$, is

$$m_{t,t+1} = \beta x_{t+1}^{-\alpha} \left(\frac{S_{t+1}}{S_t}\right)^{-\alpha} \tag{1}$$

where β is the utility discount factor, and the curvature parameter α, together with the surplus ratio, determines the agent's local degree of risk aversion.[3] As Campbell and Cochrane note, countercyclical risk aversion is a key feature of their specification.

Consumption growth x_{t+1} is assumed to be *i.i.d.* lognormal, and the log surplus ratio is assumed to evolve according to

$$\log(S_{t+1}) = (1-\phi)\bar{s} + \phi \log(S_t) + \lambda(S_t)[\log(x_{t+1}) - g] \tag{2}$$

where g is the mean of log consumption growth, ϕ controls the persistence of the surplus ratio process, and the crucial function $\lambda(S_t)$ controls the sensitivity of changes in the surplus ratio to shocks to consumption growth.[4]

The key to their reverse-engineering is the form of $\lambda(S_t)$,which is decreasing in S_t, hence countercyclical.

1.2. Melino and Yang

Melino and Yang, in their 2003 paper in the *Review of Economic Dynamics*, perform another type of reverse engineering exercise. Using Mehra and Prescott's two-state Markov chain for consumption growth, and assuming that consumption growth is a sufficient statistic for the riskless rate and the price-dividend ratio of an aggregate consumption claim, they derived the stochastic discount factor that, in combination with the Mehra–Prescott consumption process, yields equity and riskless return processes that exactly match the means and standard deviations calculated by Mehra and Prescott from their long sample of asset returns.

Recall that the Mehra–Prescott Markov chain has

$$x_t \in \{x_L, x_H\} = \{0.982, 1.054\} \tag{3}$$

and

$$P = \begin{bmatrix} P_{LL} & P_{LH} \\ P_{HL} & P_{HH} \end{bmatrix} = \begin{bmatrix} 0.43 & 0.57 \\ 0.57 & 0.43 \end{bmatrix} \tag{4}$$

where $P_{ij} = \Pr\{x_{t+1} = x_j : x_t = x_i\}$. Here, L and H denote the low and high consumption growth states, respectively. Mehra and Prescott's long sample of data on returns has an average risk-free rate of 0.8% and an average equity return of 7%. The standard deviations of the risk-free rate and equity return are 5.6 percentage points and 16.5 percentage points, respectively.

The Melino-Yang SDF is

$$\hat{m} = \begin{bmatrix} \hat{m}_{LL} & \hat{m}_{LH} \\ \hat{m}_{HL} & \hat{m}_{HH} \end{bmatrix} = \begin{bmatrix} 1.86 & 0.24 \\ 1.13 & 0.95 \end{bmatrix} \tag{5}$$

---

[1] For an example, see the 'Markov Asset Pricing' section in John Stachurski and Thomas Sargent's online lectures on quantitative economics: http://quant-econ.net/py/markov_asset.html.

[2] Models with Epstein-Zin preferences, rare disasters, concerns for robustness, disappointment aversion, and (with some effort) long-run consumption risk can all be treated computationally as simple extensions of the Mehra and Prescott's framework.

[3] Campbell and Cochrane show that, locally, relative risk aversion is given by $\left|\frac{u''(c_t-h_t)c_t}{u'(c_t-h_t)}\right| = \frac{\alpha}{S_t}$. Here and below, our notation differs slightly from Campbell and Cochrane's.

[4] Campbell and Cochrane write λ as a function of $\log(S_t)$, but that difference is immaterial here.

Any model that reproduces the SDF $\hat{m}$ within Mehra and Prescott's Markov chain framework will exactly match the first and second moments of Mehra and Prescott's asset returns data.

More suggestively, one can use $\hat{m}$ and the Mehra–Prescott Markov transition matrix (4) to derive risk-neutral probabilities, an insight of Routledge and Zin (2010). These are given by

$$\hat{\psi} = \begin{bmatrix} \hat{\psi}_{LL} & \hat{\psi}_{LH} \\ \hat{\psi}_{HL} & \hat{\psi}_{HH} \end{bmatrix} = \begin{bmatrix} 0.85 & 0.15 \\ 0.61 & 0.39 \end{bmatrix} \tag{6}$$

The Melino–Yang risk neutral probabilities suggest that countercyclical risk aversion is an important element of any resolution of Mehra and Prescott's puzzle. Conditional on being in the low-growth state, for example, the objective probability of remaining in the low-growth state is just 0.43, versus $\hat{\psi}_{LL} = 0.85$. Conditional on being in the high-growth state, the risk neutral probabilities are quite close to the objective probabilities—$\hat{\psi}_{H,j} = \{0.61, 0.39\}$ versus $P_{H,j} = \{0.57, 0.43\}$.

However, as Melino and Yang demonstrate—by trying to calibrate various standard and state-dependent preference specifications so as to produce SDFs that match $\hat{m}$—countercyclical risk aversion, while important, is alone not sufficient to resolve the puzzle.

One preference specification that Melino and Yang examine only cursorily is that of Campbell and Cochrane, as it appears to require expanding the model's state space. Melino and Yang's calculations show that the asset return data can be rationalized (using state-dependent preferences) without adding extra states.

2. A Markov-chain version of Campbell and Cochrane

Consider the log growth rate of the surplus ratio from (2),

$$\log(S_{t+1}/S_t) = (\phi - 1)[\log(S_t) - \bar{s}] + \lambda(S_t)[\log(x_{t+1}) - g]. \tag{7}$$

A key feature of Campbell and Cochrane's model is the non-constant response of growth in the surplus ratio to innovations to consumption growth, captured in the function $\lambda(S_t)$. While Campbell and Cochrane assume ϕ is close to unity, the conditional mean of $\log(S_{t+1}/S_t)$ is nevertheless non-constant as well.

Could we put Campbell and Cochrane in the Mehra–Prescott framework simply by writing the surplus ratio S_t as a function of the current Markov state? As Melino and Yang point out, that approach would not allow us to match the returns data: we'd be effectively adding only one parameter to the SDF, in addition to α and β, and our SDF would lack the flexibility necessary to match $\hat{m}$. To see this, note that across the HH or LL transitions, we would have $S_{t+1}/S_t = 1$, while the growth rates across the LH and HL transitions would be inversely related. The SDF that results would have the form

$$m = \begin{bmatrix} m_{LL} & m_{LH} \\ m_{HL} & m_{HH} \end{bmatrix} = \begin{bmatrix} \beta x_L^{-\alpha} & \theta\beta x_H^{-\alpha} \\ \theta^{-1}\beta x_L^{-\alpha} & \beta x_H^{-\alpha} \end{bmatrix}$$

where $\theta = S_H/S_L$.

Melino and Yang suggest introducing S_t as an independent state, but view this as inferior to their own state-dependent-preferences approach, which resolves the equity premium puzzle without expanding the set of Markov states.

Distinct from either of those approaches—writing S_t as a function of x_t or making S_t an additional state variable—we may note that the *level* of the surplus ratio doesn't matter for asset pricing, since the SDF depends only on the growth rate.

With that in mind, we can capture the spirit of Campbell and Cochrane's dynamics—as given in (7)—by writing the log growth rate of the surplus ratio from t to $t+1$ as a function of realized growth (x_{t+1}), with parameters that depend on the current Markov state (x_t):

$$\log(S_{t+1}/S_t) = v(x_t) + \lambda(x_t)\log(x_{t+1}). \tag{8}$$

As long as $\lambda(x)$ is non-constant, (8), like Campbell and Cochrane's (7), features time-varying conditional volatility. In Campbell and Cochrane, S_t is positively related to x_t, so we would expect $\lambda(x_t)$ to be decreasing in x_t, just as Campbell and Cochrane's reverse engineering leads them to require that their $\lambda(S_t)$ to be decreasing in S_t. This is in fact what we derive below.

Using (8), we can write the SDF (1) as

$$\begin{aligned} m_{t+1} &= \beta x_{t+1}^{-\alpha}\left(e^{v(x_t)}x_{t+1}^{\lambda(x_t)}\right)^{-\alpha} \\ &= \beta e^{-\alpha v(x_t)} x_{t+1}^{-\alpha(1+\lambda(x_t))} \end{aligned}$$

Since x_t follows a Markov chain, we can write v_i for $v(x_i)$ and λ_i for $\lambda(x_i)$, for $i = L, H$. Then, the SDF becomes

$$m = [m_{i,j}] = \left[\beta e^{-\alpha v_i} x_j^{-\alpha(1+\lambda_i)}\right]. \tag{9}$$

One drawback of this formulation—though not from the limited perspective of asset pricing—is that it renders the surplus ratio itself nonstationary. In Campbell and Cochrane's model, the surplus ratio is a stationary, though highly persistent, stochastic process. As we'll see below, though, the surplus ratio process described by (8) can be calibrated to be driftless, without affecting its ability to match the asset returns data.[5]

3. Meeting Melino and Yang

To reverse engineer the surplus ratio in the Mehra–Prescott framework, we attempt to match the SDF (9) to the Melino–Yang SDF $\hat{m}$ for a suitable choice of parameters. In other words, the problem is to find $\alpha, \beta, \{v_L, v_H\}, \{\lambda_L, \lambda_H\}$ such that

$$\beta e^{-\alpha v_i} x_j^{-\alpha(1+\lambda_i)} = \hat{m}_{i,j}, \tag{10}$$

where $\hat{m}$ is given by (5). As it turns out, there are enough parameters to accomplish this matching—for any α and β, we can find $\{v_L, v_H\}, \{\lambda_L, \lambda_H\}$ such that (10) holds.

To see this, take logs and rearrange to get

$$v_i + \log(x_j)\lambda_i = \frac{1}{\alpha}\left[\log(\beta) - \alpha\log(x_j) - \log(\hat{m}_{i,j})\right] \tag{11}$$

There are thus two equations to solve for (v_L, λ_L) and two equations to solve for (v_H, λ_H). For (v_L, λ_L), we have

$$\begin{bmatrix} 1 & \log(x_L) \\ 1 & \log(x_H) \end{bmatrix} \begin{bmatrix} v_L \\ \lambda_L \end{bmatrix} = \frac{1}{\alpha}\begin{bmatrix} \log(\beta) - \alpha\log(x_L) - \log(\hat{m}_{L,L}) \\ \log(\beta) - \alpha\log(x_H) - \log(\hat{m}_{L,H}) \end{bmatrix} \tag{12}$$

An analogous expression obtains for (v_H, λ_H). The values are determined uniquely since

$$\begin{bmatrix} 1 & \log(x_L) \\ 1 & \log(x_H) \end{bmatrix}$$

is invertible for x_L and x_H as given in the Mehra–Prescott process.

Solving (12) for $\lambda_i, i = L, H$, gives

$$\lambda_i = -1 + \frac{\log\left(\hat{m}_{i,L}/\hat{m}_{i,H}\right)}{\alpha\log(x_H/x_L)} \tag{13}$$

With the Mehra–Prescott process, $\log(x_H/x_L) = 0.0708$, approximately 2 times the standard deviation of x. For the Melino–Yang SDF, given in (5),

$$\log(\hat{m}_{L,L}/\hat{m}_{L,H}) = 2.03$$
$$\log(\hat{m}_{H,L}/\hat{m}_{H,H}) = 0.17$$

Substituting these numbers into (13) gives

$$\lambda_L = -1 + \frac{1}{\alpha}28.73$$
$$\lambda_H = -1 + \frac{1}{\alpha}2.42 \tag{14}$$

Thus, λ is strongly decreasing from the low- to high-growth state, just as Campbell–Cochrane's λ is strongly decreasing in the current surplus ratio.

For $\alpha = 1$, say, the range of our λ is in fact close to the typical range of Campbell–Cochrane's λ, if we take as typical the image under their $\lambda(S)$ of a (conditional) two standard deviation interval around their $\bar{s}$, using the law of motion (7) and their parameter values.[6] Using their parameters—from their Table 1—at an annual frequency, I calculate this range to be $[2.41, 22.80]$.

The solution for $v_i, i = L, H$, is

$$v_i = \frac{\log(\beta)}{\alpha} - \frac{1}{\alpha}\frac{\log(x_H)\log(\hat{m}_{i,L}) - \log(x_L)\log(\hat{m}_{i,H})}{\log(x_H/x_L)} \tag{15}$$

Using the values for $\log(x)$, from (3), and for $\log(\hat{m})$, from (5), gives

[5] In particular, the process can be calibrated so that $\mathbb{E}[\log(S_{t+1}/S_t)] = 0$. Absent consumption growth fluctuations from any date $t$on, our log surplus ratio would remain constant; in Campbell and Cochrane's model, it would eventually converge to its steady state level. The implicit evolution of the habit stock in our formulation is no doubt more complicated than that implicit in Campbell and Cochrane's model, though neither is readily describable by the linear processes common to other habit models.

[6] In their notation, this image is $\lambda(\bar{s} \pm 2\lambda(\bar{s})\sigma)$.

Table 1
Effect of varying consumption autocorrelation. P_{ii} and P_{ij} are the diagonal and off-diagonal elements of the Markov matrix for consumption growth; $\hat{\psi}$ is the set of risk-neutral probabilities; the α_i are the state-dependent curvature parameters; the β_i are the state-dependent discount factors; $\mathbb{E}(\beta)$ is the unconditional mean of the β_i.

$\rho =$	−0.14	0.0	0.30	0.42
P_{ii}	0.43	0.5	0.65	0.71
P_{ij}	0.57	0.5	0.35	0.29
$\hat{\psi}_{LL}$	0.852	0.862	0.886	0.897
$\hat{\psi}_{LH}$	0.148	0.138	0.114	0.103
$\hat{\psi}_{HL}$	0.611	0.582	0.510	0.475
$\hat{\psi}_{HH}$	0.389	0.418	0.490	0.525
α_L	28.727	25.857	20.206	17.978
α_H	2.416	4.698	9.306	11.257
β_L	1.105	1.013	0.887	0.857
β_H	1.078	1.123	1.292	1.403
$\mathbb{E}(\beta)$	1.092	1.068	1.090	1.130

$$\begin{aligned} v_L &= \frac{1}{\alpha}(\log(\beta) - 0.100) \\ v_H &= \frac{1}{\alpha}(\log(\beta) - 0.075) \end{aligned} \tag{16}$$

As one might expect, based on Campbell and Cochrane's calibration of their surplus process, the cyclical variation in v is much smaller than the variation in λ.

For given α and β, and using Mehra and Prescott's Markov transition matrix, (14) and (16) imply that the conditional mean of the log growth rate of the surplus ratio is

$$\begin{aligned} \mathbb{E}_L[\log(S_{t+1}/S_t)] &= \frac{\log(\beta) + 0.537}{\alpha} - 0.022 \\ \mathbb{E}_H[\log(S_{t+1}/S_t)] &= \frac{\log(\beta) - 0.046}{\alpha} - 0.012 \end{aligned} \tag{17}$$

For example, for $\alpha = 1$,

$$\begin{aligned} \mathbb{E}_L[\log(S_{t+1}/S_t)] &= \log(\beta) + 0.515 \\ \mathbb{E}_H[\log(S_{t+1}/S_t)] &= \log(\beta) - 0.058 \end{aligned}$$

As long as β is not too small, in the low-growth state the surplus ratio is expected to increase, while in the high-growth state, it is expected to decline. And, if $\log(\beta) = -(1/2)(0.515 - 0.058)$, a value of β just under 0.8, the unconditional expectation of $\log(S_{t+1}/S_t)$ will be zero. In log terms, the surplus ratio will be non-stationary, but have zero drift.

In the next section, I go one step further and map the SDF (9) into an SDF of the standard time-additive, CRRA form, but with state-dependent preference parameters (relative risk aversion and discounting). This is more squarely within Melino and Yang's framework. As in their model, the preference parameters are functions only of the current Markov state, though it is worth emphasizing that the starting point was a model with a more complex dependence of the agent's preferences on the set of Markov states. The growth rate of our surplus ratio from $t$to $t+1$ depends on both $x_t$and x_{t+1}, but the effects of the current and next-period states, conveniently, can be separated in such a way that the parameters of the state-dependent equivalent representation depend only on x_t.

4. State-dependent preferences

As claimed above, one can re-interpret the preferences we've specified here as a state-dependent version of the standard time-additively-separable, constant relative risk aversion form, with variation in both the coefficient of relative risk aversion and the utility discount factor. That is, we may re-write the SDF (9) in the form

$$m_{i,j} = \beta_i x_j^{-\alpha_i} \tag{18}$$

The mapping is easily derived from (9), defining β_i and α_i by setting

$$\beta_i x_j^{-\alpha_i} = \beta e^{-\alpha v_i} x_j^{-\alpha(1+\lambda_i)} \tag{19}$$

for $i, j = L, H$. That is,

$$\alpha_i = \alpha(1+\lambda_i)$$
$$\beta_i = \beta e^{-\alpha v_i}. \tag{20}$$

Combining (20) with (14) and (16)—or directly equating the SDF in (18) with the Melino–Yang SDF $\hat{m}$—gives the values of β_i and α_i consistent with first and second moments of the asset return data:

	L	H
α_i	28.73	2.42
β_i	1.1052	1.0782

As expected, the state-dependent representation features a strongly countercyclical risk aversion coefficient, varying from roughly 2.4 in the high-growth state to nearly 30 in the low-growth state. That the representation features state-dependent risk aversion is not surprising, given Campbell and Cochrane's interpretation of their habit specification (or the Melino–Yang risk-neutral probabilities).

More surprising is the state-dependence of the utility discount factor; in this representation, the discount factor is uniformly greater than one and countercyclical (so the rate of time preference is negative and procyclical). An agent with these preferences 'upcounts' future utility in either state, the more so (more patiently) in the low-growth state. The variation is sizable: the agent's rate of time preference differs by about 0.025, or 2.5 percentage points, across states.

Upcounting on average, of course, helps match the low average risk-free rate, a fact pointed out early on by Benninga and Protopapadakis (1990). The countercyclicality of the utility discount factor, though, is at first glance puzzling. The parameters have been reverse-engineered to match Melino and Yang's SDF, and that SDF corresponds to a countercylical risk-free rate. One might have expected a lower discount factor (and higher rate of time preference) in the low-growth state.

It turns out that, without the offsetting countercyclicality of the utility discount factor, the implied risk-free rate (as well as the implied equity return) would be too countercyclical. Precisely, suppose that we replace $\beta = \{\beta_L, \beta_H\}$ with the average of β_L and β_H (keeping the behavior of α_i the same). The resulting SDF would (roughly) match the mean risk-free rate (0.8%), but with too high a standard deviation. The model would fail on other dimensions as well.[7]

Does Campbell and Cochrane's own model—rather than just our version of it—have a state-dependent representation with a utility discount factor that's countercyclical and greater than one? Fig. 1 shows the result of simulating Campbell and Cochrane's model, at an annual frequency, using the parameters given in their Table 1. In constructing the figure, I simulated the behavior of their stochastic discount factor (for a given path of consumption growth innovations) and defined β_t by

$$\beta_t x_{t+1}^{-\alpha_t} = m_{t,t+1} \tag{21}$$

where $m_{t,t+1}$ is the realization of the SDF from period t to $t+1, x_{t+1}$ is (gross) consumption growth from t to $t+1$, and $\alpha_t = \alpha(1+\lambda(S_t))$.

The resulting β_t—simulated for 100 periods—is almost always greater than one. The lower panel of the figure plots the dependence of β_t on the log surplus ratio, verifying the countercyclicality of the utility discount factor.[8]

Melino and Yang do not consider exactly the case of an SDF given by (18); their framework features Epstein-Zin preferences, with potential variation in one or more of that family's three parameters (risk aversion, intertemporal substitution, and discounting). They do, however, look at the case of cyclical risk aversion and discounting, holding fixed the elasticity of intertemporal substitution. While that case can be calibrated to match the SDF $\hat{m}$, they rule it out, for a variety of technical reasons, on the grounds that the discount factor turns out to exceed one in one or both of the Markov states.[9]

5. A robustness check

It is natural to ask whether our result on the behavior of the utility discount factor hinges on the negative autocorrelation in consumption growth that characterizes the Mehra–Prescott Markov chain. After all, the autocorrelation properties of consumption growth can differ depending on the sample period and also with corrections or adjustments to data collected prior to the creation of systematic national accounts. Azeredo (2014), for example, estimates an autocorrelation in annual consumption growth of 0.42 for the period 1899–2012, using corrected pre-1929 data. Samples starting after 1930 also feature mild positive autocorrelation, on the order of 0.3.

Campbell and Cochrane, moreover, assume that consumption growth is i.i.d, making the zero autocorrelation case of interest as well.

7 The standard deviation of the implied risk-free rate in this case is 1.2 percentage points too high. The volatility of the implied equity return is too high by a similar magnitude, and the implied equity premium is too high by two percentage points.

8 All the MATLAB code for this paper can be found at http://www.jimdolmas.net/economics/current-work.

9 And the cyclicality they find is in fact the opposite of what we obtain here.

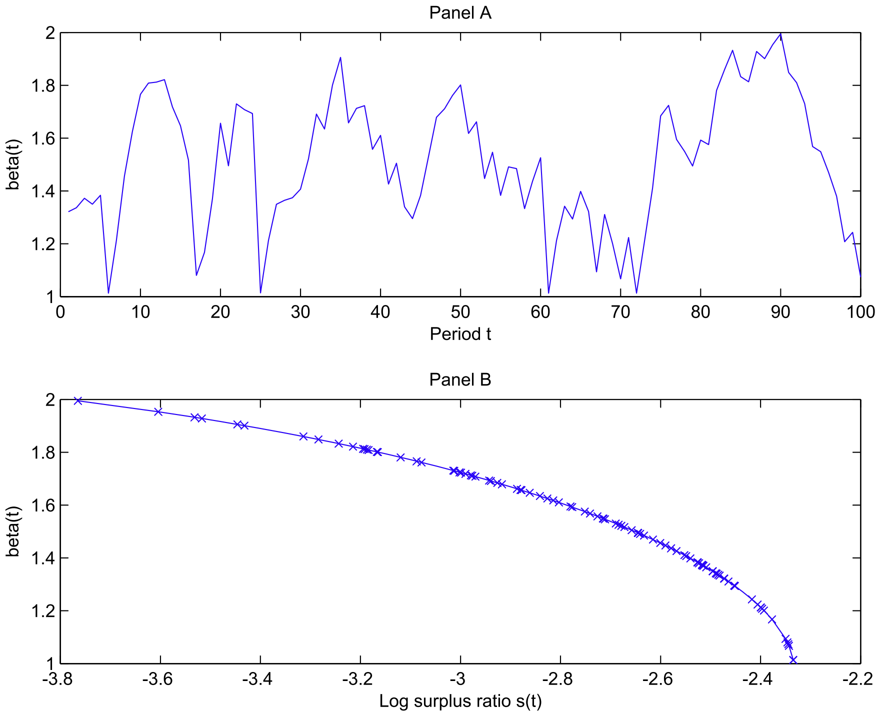

Fig. 1. State-dependent utility discount factor in the Campbell–Cochrane model. The upper panel shows β_t over time; the lower panel plots β_t against s_t, the log surplus ratio. Data are simulated using the annual versions of parameters given in Campbell and Cochrane's Table 1. The simulation starts from $s_0 = \bar{s}$, and the first 100 periods have been discarded. Consumption growth innovations were generated using MATLAB's `randn` function.

Changes to the assumed autocorrelation of consumption growth impact the calculations presented above by changing the 'reverse-engineered' stochastic discount factor $\hat{m}$ from that given in (5). Our derivations of the λ_i and v_i parameters of the surplus ratio, or the state-dependent α_i and β_i, depend on the autocorrelation in the consumption process only through $\hat{m}$.[10]

Let ρ denote the autocorrelation in consumption. To evaluate the impact of alternative choices of ρ we follow the methodology of Melino and Yang to derive alternative stochastic discount factors $\hat{m}$ that are consistent with the first and second moments of asset returns. This involves solving for the price-dividend ratios for the two consumption growth states (call them w_L and w_H) that—together with the Markov process for consumption growth—produce a matrix of equity returns across state transitions that match the unconditional mean (7%) and standard deviation (16.5 percentage points) in the data. Given the matrix of returns across state transitions and the values of the risk-free rate in the $L$and Hstates, it is straightforward to back out a matrix of risk-neutral probabilities, and from that a stochastic discount factor.[11]

The results of this exercise, for values of ρ equal to -0.14 (our benchmark case), 0.0(Campbell and Cochrane's assumption), 0.3 and 0.42 are shown in Table 1. For each value of ρ, the table records the diagonal and off-diagonal elements of the symmetric state-transition probability matrix (P) that obtains, the elements of the matrix of risk-neutral probabilities ($\hat{\psi}$) derived in the manner of Melino and Yang, and the state-dependent preference parameters that obtain by equating $[\beta_i x_j^{-\alpha_i}]_{i,j=L,H}$ to the stochastic discount factor $\hat{m}$ associated with $\hat{\psi}$, as described in Section 4. The form of the surplus ratio parameters λ_i and v_i can be backed out using the equivalence described in Section 4.

Some interesting features emerge as the autocorrelation in consumption growth is assumed to be increasingly positive. For one, the extent of countercyclicality in risk aversion needed to rationalize the returns data falls as ρ increases. This is evident first in the diminishing difference between the "objective" probabilities P_{ii} and P_{ij} and their risk-neutral counterparts, conditional on being in the Lstate, and the increasing difference conditional on being in the Hstate. The difference between α_L and α_H decreases, with the former falling by about 11and the latter increasing by about 9.

The cyclicality in β_i changes, with β_i becoming procyclical, and—for $\rho = 0.3$ or 0.42—we have upcounting of future utility only in the Hstate.[12] The upcounting in the Hstate is more extreme at higher values of ρ, and, notably, the average value of β_i remains greater than one.

[10] There is a direct dependence on the consumption growth rates in the $L$and Hstates, but these are unaffected by the choice of auotcorrelation.

[11] As in Melino and Yang, the moment conditions we are trying to match result in a quadratic equation in the price-dividend ratios w_i. There are two sets of roots, but one set is easily ruled out as violating absence of arbitrage.

[12] The unconditional volatility of β_i also increases substantially; the unconditional standard deviation (not shown) rises from about 1% of $\mathbb{E}(\beta)$ to about 24%.

6. Implications for additional moments

While Melino and Yang focus on exactly matching the first and second moments of equity returns and the risk-free rate calculated by Mehra and Prescott (and assuming Mehra and Prescott's Markov chain for consumption growth), Campbell and Cochrane compare their model's results to additional features of asset market data, including the volatility of the price/dividend ratio and the forecastability of returns.

In our framework, any additional properties of prices and returns are encoded in the stochastic discount factor $\hat{m}$, which is in turn pinned down by the first two moments of returns and the process for consumption growth.[13] Our exercise is concerned with preference specifications that rationalize $\hat{m}$. The state-dependent preferences derived above, their Campbell–Cochrane-like equivalent, and in fact any specification that, together with the process for consumption growth, matches $\hat{m}$ will have the same asset-pricing implications. The ability or failure to match additional asset-pricing moments are features of $\hat{m}$ and the Markov chain for consumption growth, regardless of the preferences that rationalize $\hat{m}$.

That said, it is worth pointing out some additional implications of $\hat{m}$ and the assumed process for consumption growth. First, since we are in a Markov chain environment, the autocorrelation of any object that depends only on the current state will be dictated by the autocorrelation inherent in the transition matrix P–i.e., the assumed autocorrelation of consumption growth. Likewise, any object that depends only on the current state will have a contemporaneous correlation of one with consumption growth. This is the case for price-dividend ratios and the risk-free rate.

Realized equity returns–as well as excess returns–depend on both the current and next-period state and so in principle may co-move less closely with consumption growth and display autocorrelations that deviate from those inherent in the transition matrix P. In practice, the model's realized returns are essentially i.i.d. for the values of ρ in Table 1. The contemporaneous correlation between excess returns and consumption growth is less than one, but still high, ranging from 0.83 at $\rho = -0.14$ to 0.67 at $\rho = 0.42$.

Volatility of the price-dividend ratio is too low at all values of ρ, but $\hat{m}$ can deliver a realistic coefficient in a predictability regression of one-period-ahead excess returns on the log dividend-price ratio.[14] For values of ρ between 0.3 and 0.4, the model-implied regression coefficients are 0.16 and 0.06, which are in the range of values found, for example in Campbell and Yogo (2006). For ρ or less, excess returns are actually *too* sensitive to the log dividend-price ratio—we obtain too much in the way of predictability.

7. Conclusion

This paper has laid out a Markov-chain-friendly version of Campbell–Cochrane preferences and from that derived a state-dependent preference equivalent. In addition to providing some pedagogical value, the exercise hopefully sheds some light on that workhorse habit model and expands the set of plausible state-dependent preference specifications. In particular, it is hard to reject out of hand state-dependent specifications with upcounting of future utility if one also accepts the plausibility of Campbell–Cochrane preferences.

Our two-state version of Campbell–Cochrane preferences maps neatly into a state-dependent preference specification where both the coefficient of relative risk aversion and the utility discount factor vary with the state. Choosing the parameters of either to exactly match first and second moments of asset returns data reveals an important role for variation in discounting, in addition to the expected countercyclicality of risk aversion. Notably, the state-dependent utility discount factor displays substantial upcounting of future utility and substantial variation across the two states.

Acknowledgment

I am grateful for the comments of an anonymous referee on an earlier version of this paper.

References

Azeredo, F. (2014). The equity premium: A deeper puzzle. *Annals of Finance, 10*(3), 347–373.
Benninga, S., & Protopapadakis, A. (1990). Time preference, leverage, and the 'equity premium puzzle'. *Journal of Monetary Economics, 25*(1), 49–58.
Campbell, J. Y., & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy, 107*(2), 205–251.
Campbell, J. Y., & Yogo, M. (2006). Efficient tests of stock return predictability. *Journal of Financial Economics, 81*(1), 27–60.
Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics, 15*, 145–161.
Melino, A., & Yang, A. X. (2003). State-dependent preferences can explain the equity premium puzzle. *Review of Economic Dynamics, 6*(4), 806–830.
Routledge, B. R., & Zin, S. E. (2010). Generalized disappointment aversion and asset prices. *Journal of Finance, 65*(4), 1303–1332.

[13] Since $\hat{m}$ depends on the specification of ρ, as shown in Section 5, more precise notation would write $\hat{m}_\rho$.
[14] Standard deviations of the log price-dividend ratio range from 0.085 to 0.124 as ρranges from -0.14 to 0.42