

REGIME LEARNING AND ASSET PRICES IN A LONG-RUN MODEL: THEORY

BIN BIN DENG*

Department of Economics
University of Chicago

November 23, 2014

PRELIMINARY
(COMMENTS ARE WELCOME)

*University of Chicago Department of Economics. Correspondence: 345 E Wacker Dr, 1901, Chicago, IL 60601. Tel: +1 202-999-6000. Email: dbb@uchicago.edu. I thank Lars Peter Hansen, Pietro Veronesi, and Xiao Qiao for comments and discussion.

Abstract

This paper tries to draw on the relative merits of both the jump risk models and the long-run risk models with a linkage established by Bayesian learning, in an attempt to improve both asset pricing approaches in producing a better mechanism for understanding asset prices regularities. Rather than treating event risk as direct jumps in the level of aggregate income, we model it as changes in the underlying state of the world, the economic regimes, which affect aggregate consumption and dividend flows through their growth and volatility's dependence on the state. Realistically, information about the state transition is imperfect in this representative agent endowment economy and agents with recursive utility perform Bayesian learning to form and update beliefs about the conditional state arrival in order to make optimal long-run consumption-investment decisions. This new learning component to the consumption-based paradigm will generate novel pricing implications through inducing extra covariance to be priced. Specifically, besides the aggregate uncertainty stemming from jump risk exposure, the presence of imperfect learning behavior also generates individual ambiguity. We shall see that such dual channels can help better explain some asset pricing regularities observed, e.g. the dual puzzles, predictability issues, time-varying conditional moments, etc., and shed some new light on the long-run cash flow news approach in asset pricing.

KEYWORDS: Equilibrium asset pricing, Recursive preferences, Long-run model, Jump risk, Markov regimes, Imperfect information, Bayesian learning

1 INTRODUCTION

From standard event risk models like Rietz (1988) and Barro (2006), we know that jump risk with fixed arrival rates will have a direct impact on equilibrium risk pricing. In this paper, we want to study the pricing implications from learning behavior by agents about the transition of the state of the world. We model events (possibly disasters) as different underlying states of the world, and transitions between the states introduce structural changes to the aggregate consumption process and thus induce jumps in asset prices. Agents learn about the state transition intensities and such learning behavior would create a positive covariance between the current realization of economic events and the future expected arrivals of them. Agents then become more optimistic/pessimistic about the future the instant a good/bad state realizes. The existence of learning behavior of agents would inevitably induce time-varying jump risk manifested as changing beliefs about the state arrivals, which should in equilibrium feed into the aggregate quantities and affect asset prices. The resulting positive covariance would generate extra learning-induced concern about jump risk in asset markets, and thus extra learning-induced risk premia, which provides a novel channel to possibly better understand some of the key asset pricing puzzles. Overall, besides the aggregate uncertainty stemming from jump risk exposure generated by structural changes in aggregate consumption, the presence of learning about the underlying state of the world generates individual ambiguity coming from imperfect learning behavior, where such dual channels might help better explain some of the observed asset pricing regularities.

The direct motivation of this paper is from the advantage of structural change or regime switching framework over direct jump framework for the aggregate consumption process in modeling jump risk. First and foremost, in standard event risk models, rare events or disasters manifested as financial market crashes are fundamentally linked to large drops in aggregate output/consumption, reflecting the huge impact of movements in real quantities on financial quantities, in which case the real effects get priced. However, we see in the data that real quantities like aggregate consumption are rather smooth over time, exhibiting far less jumps than did financial quantities like stock prices. To account for this, we propose a regime switch story in which the optimizing agents are able to maintain a smooth consumption path while their wealth-consumption ratios jump over time following stochastic transitions in the underlying economic state of the world, a property shared by the aggregate stock price-dividend ratio. This mechanism helps capture the fact that we do observe a smooth sample path of aggregate consumption even in the presence of financial market crashes in reality. In this way, a smooth aggregate consumption path and volatile asset prices can be more coherent. Next, standard event risk models usually assume that the large drops in aggregate consumption upon a bad state hits realize themselves over a single period of time, which as a result requires unrealistically large one-time shocks to the aggregate income process. Even though some authors (e.g. Barro, 2006) allow for the degrees of freedom to calibrate such one-time crash length, it is still more of an ad hoc choice of parameters upon a leap of faith for the reality from the very beginning. In

contrast, by modeling jump risk as stochastic changes in some fundamental economic state of the world, large consumption drops can now realize themselves over longer horizons, say several years or even decades, without restriction of which as a condensed impact, which is more consistent with the data observed.¹

In fact, the economic framework herein is essentially similar to the long-run risk models, e.g. Bansal and Yaron (2004) (BY), in that both pursue cash flow news for equilibrium pricing and both emphasize the separation of elasticity of inter-temporal substitution (EIS) from risk aversion. Yet, differences are apparent. First, while the level of aggregate consumption has continuous sample paths here, both its growth and volatility are subject to jumps that capture disasters and recoveries. On the contrary, consumption moments in the long-run risk models remain continuous processes. Second, the key pricing channel in the long-run risk models is the persistent changes in consumption moments, which have far stronger impact on asset prices than contemporaneous consumption shocks. In contrast, aggregate consumption in this paper follows a random diffusion process conditional on a given state of the world. Notably, one of the major criticisms against the long-run risk models is that they rely strongly on the imposed high persistence of the aggregate consumption growth process, suggesting strong predictability of growth by valuation ratios, like the price-dividend ratio, which is not true in the data. This paper tries to tackle this issue by introducing a learning channel that would produce endogenous uncertainty persistence, which in turn would induce jumps in the valuation ratios, while keeping a consumption process with continuous sample paths and with no requirement for a persistent component in the growth process coming from the state evolution.² We emphasize that instead of generating the necessary persistence from the aggregate consumption growth process for pricing in such a cash flow news approach, we will see later that the proposed Bayesian learning channel serves the same role yet by generating imperfect learning-induced uncertainty persistence. Given that we lack strong empirical evidence in support of persistent fluctuations in aggregate consumption growth, to still try resolving asset pricing puzzles from the long-run cash flow news perspective, the learning mechanism proposed herein would seem more attractive.

Recently, several other authors have delved into the event risk literature by incorporating time-

¹ In addition, we can capture such durations empirically by calibrating the recovery rate after a bad state realization in the economy. Also, from possible extensions to the simple learning model laid out in Section 3, we can potentially see one more advantage being that the potentially generalized model allows for recoveries from a bad state to a good state again, a channel that was absent in standard event risk models where they focus only on normal-to-disaster phase transition.

²Empirically, although the results in BY match the first-order autocorrelation of annual consumption growth in the data, Beeler and Campbell (2009) (BC) point out that this match is critically dependent on the use of consumption data from the period of the Great Depression. If we were using postwar consumption data, BC find that the first-order autocorrelation would be much lower. Also, BC report that empirically, higher-order consumption autocorrelations, particularly the third and fourth, are very low compared to what would be predicted by BY. Preliminary empirical results to this theory paper suggest that such observation might be more consistent with implications from the model herein where higher-order autocorrelations are all closer to zero.

varying jump risk probabilities in different ways to try improving the empirical relevance of this approach as well as reaching a higher bar of the match of moments. Typically, Gabaix in his 2012 QJE paper makes use of the linearity-generating (LG) processes to collapse the modeling dynamics into one single state variable of the economy, i.e. the resilience of assets to shocks, so as to achieve time-varying jump risk probabilities and at the same time obtain close-form solutions to asset prices and returns. The gist of that technicality lies in the use of an LG twist forced into the state evolution process that produces a quadratic term in the drift of the aggregate income process. To a large extent, this is a result of pure reverse-engineering. Nevertheless, in our approach, we have effectively also exploited the LG twist but in a more natural way in the sense that rather than inserting a quadratic term in the mean of the linearity-generating process, learning-induced covariance naturally shows up in the diffusion part and helps kill off the quadratic term. In this sense, a learning-induced time-varying jump risk exposure modeling approach seems more attractive. In addition, Gabaix (2012) uses power period utility function for the representative agent, which implies that an increase in the disaster probability reduces the risk-free rate and asset prices go up as a result, which is rather counter-intuitive. In this paper, we use recursive preferences in the form of EZ utility, which helps resolve this undesirable feature by the separation of EIS and risk aversion.

2 MODEL ENVIRONMENT

2.1 Preferences

We follow the standard literature assuming complete markets in a representative agent economy. Agents in this economy are assumed to have continuous-time recursive preferences in the form of instantaneous Epstein-Zin (EZ) utility

$$V_t = \left\{ \left(1 - e^{-\delta\varepsilon}\right) (rC_t)^{\frac{1-\gamma}{\theta}} + e^{-\delta\varepsilon} \left[\mathbb{E}_t V_{t+\varepsilon}^{1-\gamma}\right]^{\frac{1}{\theta}} \right\}^{\frac{\theta}{1-\gamma}} \quad (1)$$

parametrized by the relative risk aversion γ , the elasticity of inter-temporal substitution $1/\rho$, the continuous discount rate δ , and the scaling parameter $r > 0$, where $\theta \equiv \frac{1-\gamma}{1-\rho}$. To emphasize, we will see the importance of using recursive preferences later. Note that we can equivalently express the preferences by the stochastic differential utility (SDU) as in Duffie and Epstein (1992).

2.2 Regime Cycle

As described in Section 1, we think of event risk (including disasters) as changing regimes determined by the underlying state of the world. For ease of illustration, suppose there are only two states of the world, good or bad, where the good state corresponds to a normal/good economic regime that

gives a higher consumption growth with a lower volatility, while the bad state corresponds to a disastrous/bad economic regime that gives the reverse. Let $\{X_t\}$ be a stationary discrete-state Markov process, realized as coordinate vectors in R^2 , i.e. $X_t = x_1 = [1, 0]'$ or $x_2 = [0, 1]'$. In this case, X_t simply picks up a growth and volatility regime at each point in time. Assuming an intensity matrix A for $\{X_t\}$, the transition probability matrix over the time interval ε is thus given by $P_{A,\varepsilon} \equiv \exp(\varepsilon A)$. When a bad state hits, the level of consumption per se does not jump directly as what was commonly modeled in the standard disaster literature, but the consumption process suffers from a structural change in both its growth and volatility, induced by their dependence on the underlying state of the world X_t .

2.3 Aggregate Consumption and Dividend

Conditional on a given state of the world at time t , we assume a geometric diffusion process for the aggregate consumption

$$d\ln C_t = \mu_c(X_t)dt + \sigma_c(X_t)dW_t \quad (2)$$

where W_t is a standard Brownian motion same to all states. To be specific, assume the form

$$d\ln C_t = (\beta_c \cdot X_t + \mu_c)dt + \sigma_c \cdot X_t dW_t \quad (3)$$

In particular, we restrict $\mu_c(x_2) < \mu_c(x_1)$ and $\sigma_c(x_2) > \sigma_c(x_1)$. Following Martin (RES 2013), we model dividends of the aggregate stock as a levered consumption process

$$D_t \equiv C_t^h$$

where the parameter h denotes a proxy for the level of leverage that scales the volatility of the aggregate stock. When $h = 0$, the asset is riskless, and when $h = 1$, the asset is the aggregate wealth portfolio that pays out aggregate consumption flows. Given the aggregate consumption process (3), by Ito's lemma, we obtain the following implied aggregate dividend process as a function of the underlying state X_t

$$d\ln D_t = \mu_D(X_t)dt + \sigma_D(X_t)dW_t \quad (4)$$

where $\mu_D(X_t) \equiv h(X_t' \beta_c + \mu_c) - \frac{1}{2}h(1-h)X_t' \sigma_c \sigma_c' X_t$ and $\sigma_D(X_t) \equiv hX_t' \sigma_c$. Let the price of aggregate stock at time t be P_t^S and denote the aggregate price-dividend (P/D) ratio as $\Phi \equiv \frac{P_t^S}{D_t}$, which is a function of X_t . Note that P_t^S is a valuation functional and its parametrization and return dynamics will be discussed in later sections.

3 SIMPLE LEARNING: MARKOV STRUCTURE ON BAD STATE ARRIVAL

In this economy, the source of uncertainty that induces learning behavior is the imperfect information about the state transition, i.e. agents do not know the true transition probabilities, partially or completely, so they have to learn:

$$P_{A,\varepsilon} \equiv \begin{bmatrix} good_t \rightarrow good_{t+\varepsilon} & good_t \rightarrow bad_{t+\varepsilon} \\ bad_t \rightarrow good_{t+\varepsilon} & bad_t \rightarrow bad_{t+\varepsilon} \end{bmatrix} \equiv \begin{bmatrix} 1 - p_{1,\varepsilon} & p_{1,\varepsilon} \\ p_{2,\varepsilon} & 1 - p_{2,\varepsilon} \end{bmatrix}, \forall t \quad (5)$$

whose stationary probability vector is $q_{A,\varepsilon} \equiv \begin{bmatrix} good_\varepsilon \\ bad_\varepsilon \end{bmatrix} = \begin{bmatrix} \frac{p_{2,\varepsilon}}{p_{1,\varepsilon} + p_{2,\varepsilon}} \\ \frac{p_{1,\varepsilon}}{p_{1,\varepsilon} + p_{2,\varepsilon}} \end{bmatrix}$ that satisfies $q'_{A,\varepsilon} = q'_{A,\varepsilon} P_{A,\varepsilon}$, with $P_{A,\varepsilon} = exp(\varepsilon A)$ and the true intensity matrix

$$A \equiv \begin{bmatrix} -\sum_{j \neq g} \lambda_{gj} & \lambda_{gb} \\ \lambda_{bg} & -\sum_{j \neq b} \lambda_{bj} \end{bmatrix} \equiv \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}$$

Here, we consider a simple case of learning for ease of illustration, in which we assume that if it were a bad state at time t , agents know perfectly how likely it is going to recover over a given time interval ε , i.e. they know exactly what $p_{2,\varepsilon}$ is. However, if it were instead a good state, agents would not be able to know how likely a bad time is going to come, i.e. they do not know $p_{1,\varepsilon}$ and thus have to partially learn about A . Therefore, the conditional state transition would look like

$$q'_{A,t+\varepsilon} = \begin{bmatrix} good_{t+\varepsilon} & bad_{t+\varepsilon} \end{bmatrix} = q'_{A,t} P_{A,t,\varepsilon} = \begin{bmatrix} good_t & bad_t \end{bmatrix} \begin{bmatrix} 1 - p_{1,t,\varepsilon} & p_{1,t,\varepsilon} \\ p_{2,\varepsilon} & 1 - p_{2,\varepsilon} \end{bmatrix} \quad (6)$$

where p_1 now depends on t and $p_{1,t,\varepsilon}$ is the conditional arrival probability of a bad state over the time interval ε , conditioning on being in a good state at time t . Essentially, we are introducing a time-varying jump risk feature into the model. Next, we turn to model the learning of A . Although agents could not fully observe the realization of A at each point in time, they know a pre-determined structure on the conditional bad state arrival intensity a . Without loss of generality, assume there are two possible values a_h and a_l ($a_h > a_l$) for the true rate. Further, suppose this true rate follows a continuous-time Markov structure given by the intensity matrix

$$B \equiv \begin{bmatrix} -\sum_{j \neq l} \lambda_{lj} & \lambda_{lh} \\ \lambda_{hl} & -\sum_{j \neq h} \lambda_{hj} \end{bmatrix} = \begin{bmatrix} -\lambda_{lh} & \lambda_{lh} \\ \lambda_{hl} & -\lambda_{hl} \end{bmatrix} \quad (7)$$

with transition probability matrix over the time interval ε given by $P_{B,\varepsilon} \equiv exp(\varepsilon B)$. Thus, even though agents still have uncertainty over A , they know perfectly that the conditional bad state

arrival intensity would switch between a_h and a_l over time according to

$$\begin{cases} Pr(a_{t+dt} = a_h | a_t = a_l) = & \lambda_{lh} \\ Pr(a_{t+dt} = a_l | a_t = a_h) = & \lambda_{hl} \end{cases} \quad (8)$$

and as a result, the stationary (long-run average) arrival intensity of a bad state is given by

$$\bar{a} \equiv \frac{\lambda_{hl}}{\lambda_{lh} + \lambda_{hl}} a_l + \frac{\lambda_{lh}}{\lambda_{lh} + \lambda_{hl}} a_h \quad (9)$$

where \bar{a} is common knowledge among all agents. Therefore, to learn about the arrival intensity of a bad state at the next instant during a good regime at time t , agents form time t conditional expected probabilities of the Markov transition on a upon observing the entire history of information up to time t as

$$\begin{cases} \alpha_t & \equiv Pr(a_t = a_l | \mathcal{F}_t) \\ 1 - \alpha_t & \equiv Pr(a_t = a_h | \mathcal{F}_t) \end{cases} \quad (10)$$

Note that these time t posteriors α_t and $1 - \alpha_t$ are the expected subjective probabilities conditional on the current state being good, with long-run averages $\bar{\alpha} = \frac{\lambda_{hl}}{\lambda_{lh} + \lambda_{hl}}$ and $1 - \bar{\alpha} = \frac{\lambda_{lh}}{\lambda_{lh} + \lambda_{hl}}$.

Remark 1. *This $\{\alpha_t\}$ process captures the learning behavior of the agents in this economy and we want to characterize its dynamics over time, i.e. how do agents adjust their posterior beliefs to information updates.*

Before fully characterizing $\{\alpha_t\}$, let us perform a thought experiment. Intuitively, we should expect that if it were a good regime at the next instant in time, agents would then be less alerted about potential future bad times and update their beliefs by assigning a lower arrival intensity to a bad state in the future. On the other hand, if it were a bad regime at the next instant in time, agents would then become more panic about potential future bad times and update their beliefs by assigning a higher arrival intensity to a bad state again in the future. Note that such effects might be asymmetric, i.e. a bad state realization at the next instant in time may induce a relatively larger adjustment in beliefs about the bad state arrival intensity than a good state realization at the next instant in time does. Now, we turn to formally characterizing the dynamics of α_t over time in order to impose a proper structure to the evolution of the state of learning in this economy.

Lemma 1. *Given the information structure described above, the posterior belief α_t evolves according to the following process:*

$$d\alpha_t = \xi_t dt + \kappa_t \alpha_t dN_t$$

where $\xi_t \equiv \alpha_t (a_t^f - a_l) + (1 - \alpha_t) \lambda_{hl} - \alpha_t \lambda_{lh}$, $a_t^f \equiv \alpha_t a_l + (1 - \alpha_t) a_h \in (a_l, a_h)$, $\kappa_t \equiv (a_l - a_t^f) / a_t^f$, and dN_t is a normalized Poisson process for the bad state arrival at the next instant in time con-

ditional on a current good state; $dN_t/dt = 1$ or 0 (arrives or not), with respect to the posterior measure that has arrival rate a_t^f , which is the posterior subjective update of a_t after time t learning.³

Given this lemma, we can now confirm the previous thought experiment. Say, we initiate the α_t process at its long-run average $\bar{\alpha}$, then we have $\xi_t = \bar{\alpha} (a_t^f - a_l) > 0$. Thus, if it were a good state at the next instant in time, i.e. $dN_t/dt = 0$, we would obtain $d\alpha_t/dt = \xi_t > 0$, i.e. we should see an upward adjustment of α_t , which means that the agents now believe the economy is less likely to go into a bad state in the near future and thus update their posterior subjective belief a_t^f downward (closer to a_l). On the other hand, if it were a bad state at the next instant in time, i.e. $dN_t/dt = 1$, then (see Appendix A1)

$$d\alpha_t/dt = \frac{\bar{\alpha}}{a_t^f} (a_t^f - a_l) (a_t^f - 1) < 0$$

i.e. we should see a downward adjustment of α_t , which means that the agents now believe the economy is more likely to go into a bad state in the near future and thus update their posterior subjective belief a_t^f upward (closer to a_h). This is exactly what our previous thought experiment would predict. From now on, we can express the time t conditional expected (subjective) state transition intensity matrix as $A_t^f \equiv \begin{bmatrix} -a_t^f & a_t^f \\ b & -b \end{bmatrix}$.

Remark 2. *In what follows, we treat α_t , and thus the posterior subjective belief update a_t^f , as the extra endogenous state of the world, and together with X_t , we define the learning-adjusted state of the world as $s_t \equiv (X_t, \alpha_t) \in ((x_1, \alpha_t), (x_2, \alpha_t)) \equiv (s_{1t}, s_{2t})$, which will be critical in equilibrium pricing.*

Imperfect Learning and Persistent Uncertainty Why agents cannot perfectly learn? If the state transition parameters are constant, then agents in this economy should be able to learn perfectly in the long run. Nevertheless, given the time-varying state transition intensities, there would be persistent uncertainty in the economy about the true arrival rate of a regime/event/disaster at any point in time. In fact, we could consider a case in which the true transition probability densities are subject to small pre-specified jumps over time, which should give us constant uncertainty throughout, which is exactly the case of our simple model above.

³This lemma directly modifies from Theorem 19.6 of Liptser and Shiryaev (2001).

4 EQUILIBRIUM PRICING

4.1 SDF Dynamics

First, rewrite (1) into a continuous-time recursive equation as follows.

$$0 = -\frac{\delta}{1-\rho} \left[\left(\frac{V_t}{C_t} \right)^{1-\rho} - r^{1-\rho} \right] + \frac{1}{1-\gamma} \left(\frac{V_t}{C_t} \right)^{\gamma-\rho} \cdot \lim_{\varepsilon \downarrow 0} \frac{\mathbb{E}_t \left[\left(\frac{V_{t+\varepsilon}}{C_{t+\varepsilon}} \right)^{1-\gamma} \left(\frac{C_{t+\varepsilon}}{C_t} \right)^{1-\gamma} \right] - \left(\frac{V_t}{C_t} \right)^{1-\gamma}}{\varepsilon} \quad (11)$$

To solve for the stochastic discount factor (SDF), we conjecture the valuation-consumption (V/C) ratio as a function of the underlying learning-adjusted state: $\frac{V_t}{C_t} \equiv \phi(s_t)$ and let $\phi(s_t)^{1-\gamma} \equiv \psi(s_t)$.

Given that the wealth-consumption (W/C) ratio can be expressed as $\frac{W_t}{C_t} = (1 - \exp(-\delta\varepsilon))^{-1} \left(\frac{V_t}{C_t} \right)^{1-\rho}$ (see Appendix A2), we should have

$$\psi(s_t) = (1 - \exp(-\delta\varepsilon))^\theta \left(\frac{W_t}{C_t} \right)^\theta \equiv (1 - \exp(-\delta\varepsilon))^\theta \Lambda(s_t)^\theta \quad (12)$$

i.e. ψ can be expressed as a function of the aggregate W/C ratio denoted by Λ as a function of the underlying state s_t .

To obtain an analytical solution to the SDF, we consider two limiting cases.

Case 1: $r = 1$ and $\rho \rightarrow 1$, i.e. the log utility case. In this economy, we can easily see that the equilibrium wealth-consumption ratio is constant and we do not obtain any novel pricing implications from the learning behavior. This is intuitive because with log utility, the agents are effectively myopic and value the contemporaneous consumption so much that they essentially care little about the long-run portfolio adjustment and the consumption dynamics, and in other words, the regime learning dynamics can find no way feeding into the marginal rate of substitution between inter-temporal consumptions and thus can never show up in the SDF for equilibrium pricing.

Case 2: $r^{1-\rho} \rightarrow 0$, i.e. the long-run case. In this economy, contrary to the previous scenario, agents are made to care much more about the future than the present, i.e. long horizon consumption dynamics weigh in much more than before and thus the regime learning element with persistent uncertainty naturally finds its way feeding into the inter-temporal consideration of optimal consumption-investment decisions, which in turn affects the equilibrium asset prices. Under this parameter restriction, (11) becomes

$$\frac{\delta(1-\gamma)}{1-\rho} \left(\frac{V_t}{C_t} \right)^{1-\gamma} = \lim_{\varepsilon \downarrow 0} \frac{\mathbb{E}_t \left[\left(\frac{V_{t+\varepsilon}}{C_{t+\varepsilon}} \right)^{1-\gamma} \left(\frac{C_{t+\varepsilon}}{C_t} \right)^{1-\gamma} \right] - \left(\frac{V_t}{C_t} \right)^{1-\gamma}}{\varepsilon} \quad (13)$$

which is a linear equation in $(V_t/C_t)^{1-\gamma} = \psi(s_t)$. Appendix A3 shows that an eigenvalue problem can then be established and the critical function ψ of the state can be solved for.

Remark 3. *From now on, we work under Case 2 in a long-run model for equilibrium pricing.*

Next, note that under recursive utility, the instantaneous SDF satisfies

$$\begin{aligned} \frac{S_{t+\varepsilon}}{S_t} &= e^{-\delta\varepsilon} \left(\frac{C_{t+\varepsilon}}{C_t} \right)^{-\rho} \left(\frac{V_{t+\varepsilon}^{1-\gamma}}{\mathbb{E}[V_{t+\varepsilon}^{1-\gamma}|\mathcal{F}_t]} \right)^{\frac{\rho-\gamma}{1-\gamma}} \\ \Rightarrow S_t &= e^{-\delta t} \left(\frac{C_t}{C_0} \right)^{-\rho} \left(\frac{V_t^{1-\gamma}}{\mathbb{E}[V_t^{1-\gamma}|\mathcal{F}_0]} \right)^{\frac{\rho-\gamma}{1-\gamma}} \\ &= e^{-\delta t} \left(\frac{C_t}{C_0} \right)^{-\rho} (V_t^*)^{\frac{\rho-\gamma}{1-\gamma}} \end{aligned}$$

where $V_t^{1-\gamma} \equiv V_t^* \hat{V}_t$ and $\hat{V}_t \equiv \mathbb{E}[V_t^{1-\gamma}|\mathcal{F}_0]$. Therefore, the SDF dynamics read

$$\begin{aligned} d\ln S_t &= -\delta dt - \rho d\ln C_t + \frac{\rho-\gamma}{1-\gamma} d\ln V_t^* \\ &= \left(-\delta - \rho X_t' \beta_c - \rho \mu_c \right) dt - \rho X_t' \sigma_c dW_t + \frac{\rho-\gamma}{1-\gamma} d\ln V_t^* \end{aligned} \quad (14)$$

Notice that $d\ln V_t^*$ is by construction a multiplicative martingale and its risk exposure comes only from two sources

1. the diffusion exposure from $(1-\gamma)d\ln C_t$ that reads $\xi_v^*(X_t) \equiv \xi_v^* \cdot X_t = (1-\gamma)X_t' \sigma_c$, and
2. the jump exposure from $d\ln \psi(s_t)$ that reads $\chi_v^*(s_{t+}, s_t) \equiv \ln \psi(s_{t+}) - \ln \psi(s_t) = \ln \frac{\psi(s_{t+})}{\psi(s_t)}$ where $\chi_v^*(s, s) = 0$.

Now, suppose we guess the deterministic drift component of $d\ln V_t^*$ is of the form $\beta_v^* \cdot X_t + \mu_v^*$, then derivation in Appendix A4 shows that

$$\begin{aligned} d\ln S_t &= \beta_s(s_t) dt + \xi_s(s_t) dW_t + \chi_s(s_{t+}, s_t) \quad (15) \\ \text{where } \beta_s &= -\delta - \rho X_t' \beta_c - \rho \mu_c + \frac{\rho-\gamma}{1-\gamma} \left(X_t' \beta_v^* + \mu_v^* \right) \\ \xi_s &= -\gamma X_t' \sigma_c \\ \chi_s(s_{t+}, s_t) &= \ln \left(\frac{\psi(s_{t+})}{\psi(s_t)} \right)^{\frac{\theta-1}{\theta}} \\ X_t' \beta_v^* + \mu_v^* &= \begin{cases} x_1' \beta_v^* + \mu_v^* &= a_t^f - a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} - \frac{1}{2} (1-\gamma)^2 x_1' \sigma_c \sigma_c' x_1, \quad X_t = x_1 \\ x_2' \beta_v^* + \mu_v^* &= b - b \frac{\psi(x_1, \alpha_{t+})}{\psi(x_2, \alpha_t)} - \frac{1}{2} (1-\gamma)^2 x_2' \sigma_c \sigma_c' x_2, \quad X_t = x_2 \end{cases} \quad (16) \end{aligned}$$

Now, given $\beta_s(s_t)$, $\xi_s(s_t)$ and $\chi_s(s_{t+}, s_t)$, define $\hat{K} \equiv [\exp(\chi_s(s_{t+}, s_t))]_{2 \times 2}$, $\hat{D} \equiv \text{diag}(\beta_s) + \frac{1}{2} \text{diag}(\xi_s \xi_s')$, and $\hat{\mathbb{B}} \equiv \hat{K} \circ A^f + \hat{D}$. As a result, we can obtain the local risk-free rate in matrix form as $r^f(s) = -\hat{\mathbb{B}}1_2$. Intuitively, since there is no local risk exposure here, by the local pricing restriction we would expect the risk-free rate to be parametrized only by the SDF. Also, given the valuation functional V_t parametrized by $\beta_v(s_t)$, $\xi_v(s_t)$ and $\chi_v(s_{t+}, s_t)$, by local risk-return trade-off pricing restriction, we would have the general expression for the local expected excess return as

$$\begin{aligned} r_t^v(s_t) - r_t^f(s_t) &= -\xi_s(s_t) \cdot \xi_v(s_t) \\ &\quad - \int [1 - \exp(\chi_s(s_{t+}, s_t))] [1 - \exp(\chi_v(s_{t+}, s_t))] \eta(ds_{t+}|s_t) \end{aligned} \quad (17)$$

$$\begin{aligned} \Rightarrow r_t^e(s_t) &= -\xi_s(s_t) \cdot \xi_v(s_t) \\ &\quad - \left(X_t' A_t^f \right) \begin{bmatrix} [1 - \exp(\chi_s(s_1, s_t))] [1 - \exp(\chi_v(s_1, s_t))] \\ [1 - \exp(\chi_s(s_2, s_t))] [1 - \exp(\chi_v(s_2, s_t))] \end{bmatrix} \end{aligned} \quad (18)$$

from which we can clearly see that even though aggregate consumption still has continuous sample paths, the pricing implications for jump (regime switching) risk exposure are not degenerate.⁴ To price any cash flow processes/claims in this context, we just need to find out the parametrization of the corresponding valuation functionals, specifically $\xi_v(s_t)$ and $\chi_v(s_{t+}, s_t)$.

4.2 Aggregate Stock Price and Wealth Processes

First, recall that aggregate dividend follows a levered aggregate consumption process in this economy, and the aggregate stock is the aggregate wealth portfolio giving out consumption streams if the leverage ratio $h = 1$. Therefore, analogous to the aggregate W/C ratio Λ , the aggregate P/D ratio Φ should also be a log-linear function of the underlying state s_t and subject to the jump risk exposure coming from the learning state α_t . Next, recall from equation (4) that the aggregate dividend process is a function of the economic state X_t only. Thus, the aggregate stock price $P_t^S = D_t(X_t) \Phi(s_t)$ has local risk exposure coming from two sources

1. the diffusion exposure from $d \ln D_t$ that reads $\xi_{PS}(X_t) \equiv \xi_{PS} \cdot X_t = h X_t' \sigma_c$, and
2. the jump exposure from $d \ln \Phi(s_t)$ that reads $\chi_{PS}(s_{t+}, s_t) \equiv \ln \Phi(s_{t+}) - \ln \Phi(s_t) = \ln \frac{\Phi(s_{t+})}{\Phi(s_t)}$ where $\chi_{PS}(s, s) = 0$.

Then, by applying Ito's lemma with jumps to P_t^S , we obtain:

1. Conditional on a good state realization at time t , i.e. $X_t = x_1$, we have

$$d \ln P_t^S = \left[\mu_D(x_1) + \xi_t \frac{\partial \ln \Phi(x_1, \alpha_t)}{\partial \alpha} \right] dt + \sigma_D(x_1) dW_t + \left[\frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} - 1 \right] dN_t \quad (19)$$

⁴Note that the $\eta(\cdot|\cdot)$ here is the conditional jump density for the jump risk component.

where dN_t is a standard Poisson process governing the bad state arrival given a good state as defined in Section 3 before.

2. Conditional on a bad state realization at time t , i.e. $X_t = x_2$, we have

$$d\ln P_t^S = \mu_D(x_2) dt + \sigma_D(x_2) dW_t + \left[\frac{\Phi(x_1, \alpha_t)}{\Phi(x_2, \alpha_t)} - 1 \right] d\tilde{N}_t \quad (20)$$

which in fact does not depend on the conditional bad state arrival intensity α_t , and $d\tilde{N}_t$ is a standard Poisson process for the good state arrival given a bad state similarly defined.⁵

Note that when $h = 1$, we can similarly derive the aggregate wealth process and it should be driving us to the same results as obtained by the SDF analysis above since $\Lambda(s_t) = (1 - \exp(-\delta\varepsilon))^{-1} \psi(s_t)^{1/\theta}$.

4.3 Local Expected Excess Return

With the above parametrization of the aggregate stock price dynamics and the SDF dynamics derived under learning, we can write out the local expected excess return on the aggregate stock as follows (see Appendix A5).

1. Conditional on a good state realization at time t , i.e. $X_t = x_1$, we have

$$EP_t^S(x_1, \alpha_t) = h\gamma x_1' \sigma_c \sigma_c' x_1 - a_t^f \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \right] \quad (21)$$

where the first term corresponds to the pricing of diffusion risk and the second term corresponds to the pricing of jump risk. Notice that different from standard event risk models, the second term does not only depend on the economic state X_t , but also, and more importantly so in this model, depend on the learning state α_t that captures the agents' expectation evolution. Therefore, in this model, the covariance between the current state realization and the agents' expectation about future states given all past and contemporaneous information will play a significant role in determining the equilibrium level of risk premium. We emphasize that in this model, there is no direct jumps in the aggregate consumption level, i.e. the consumption paths of agents are still continuous, and the consumption risk manifests itself in the form of structural change in the mean and the volatility of aggregate consumption growth due to jumps in the underlying state of the world s_t . In addition, we emphasize that in order to ensure consistent pricing implications for the aggregate stock and the aggregate wealth portfolio, restrictions on the P/D as well as the W/C ratios should be considered (see Appendix A7), which could in turn be used to numerically solve for the two ratios.

⁵Note that we obtain the second equation under the previously-made assumption that there is no learning when the current state is bad.

2. Conditional on a bad state realization at time t , i.e. $X_t = x_2$, we have

$$EP_t^S(x_2, \alpha_t) = h\gamma x_2' \sigma_c \sigma_c' x_2 - b \left[1 - \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_1, \alpha_t)}{\Phi(x_2, \alpha_t)} \right] \quad (22)$$

where in this case the second term in fact does not depend on the evolution of α_t , i.e. learning happens only during a good state about the potential arrival of a bad state in the future.

ICAPM Interpretation Note that absent learning, the equity premium EP_t^S collapses to a form similar to what would have been obtained as in Campbell and Vuolteenaho (2004) (CV), i.e. $EP_t^S(x) = \rho\theta\sigma_{PS_c}(x) + (1-\theta)\sigma_{PS_w}(x)$,⁶ where the first term captures the (now time-varying) covariance of the aggregate stock with consumption growth and the second term captures the (now time-varying) covariance of the asset with the wealth portfolio held by the agents, which could be manipulated into a version of Merton's ICAPM (now with stochastic volatility) with two covariance pricing terms, i.e. $EP_t^S(x) = \gamma\sigma_{PS_w}(x) + (\gamma-1)\sigma_{PS_h}(x)$ and $\sigma_{PS_c}(x) = \sigma_{PS_w}(x) + (1-1/\rho)\sigma_{PS_h}(x)$, where the first term captures that the agents care about the covariance of the asset with their wealth portfolio, while the second term captures that they also care about the covariance with changes in their investment opportunities⁷ that are in this case caused only by the structural changes in the aggregate dividend (levered consumption) growth. Yet, in our model where persistent uncertainty about the economic state transitions exists, the endogenous learning state α_t should play a role in the ICAPM formulation and the covariance induced by the learning behavior should be priced. Intuitively, the ICAPM representation in this case should consist of three covariance pricing terms, one capturing the covariance with the aggregate consumption growth as before, one capturing the covariance with outside/exogenous changing "investment opportunities" induced by economic structural change (X_t), and the third/novel one capturing the covariance with perceived/endogenous changing "investment opportunities" induced by the Bayesian learning dynamics (α_t).⁸

4.4 Local Risk-Free Rate

Similarly to the above, we can obtain the local risk-free rate as follows (see Appendix A6).

⁶The exogenous state x follows a jump process in our model rather than some continuous process as in CV (2004).

⁷In this model, since the aggregate dividend is just levered aggregate consumption and we consider only the aggregate stock/levered wealth portfolio, we do not really have a well-defined investment opportunity set other than a changing levered consumption payout stream affecting the returns on investing in the levered wealth portfolio.

⁸If we could write out analytically the two valuation ratios as log-linear functions of the underlying states, we could in turn derive an explicit ICAPM formulation.

1. Conditional on a good state realization at time t , i.e. $X_t = x_1$, we have

$$\begin{aligned} r_t^f(x_1, \alpha_t) &= \delta + \rho \left(x_1' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 \\ &\quad + a_t^f - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} - \frac{\theta-1}{\theta} \left(a_t^f - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta} \right) \end{aligned} \quad (23)$$

where the first line corresponds to the standard expression of local risk-free rate derived in a traditional C-CAPM model if we restrict $\rho = \gamma$, i.e. the case of power utility. The second line reflects the pricing influence from both the consumption risk induced by structural change and the learning behavior of agents, which we shall illustrate separately in detail later.

2. Conditional on a bad state realization at time t , i.e. $X_t = x_2$, we have

$$\begin{aligned} r_t^f(x_2, \alpha_t) &= \delta + \rho \left(x_2' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_2' \sigma_c \sigma_c' x_2 \\ &\quad + b - b \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^{\theta-1} - \frac{\theta-1}{\theta} \left(b - b \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^{\theta} \right) \end{aligned} \quad (24)$$

where in this case, like before, the W/C ratio terms do not depend on the evolution of α_t , i.e. learning happens only during a good state about the potential arrival of a bad state in the future.⁹

Solving for $\Lambda(X_t, \alpha_t)$ and $\Phi(X_t, \alpha_t)$ To complete the equilibrium pricing of the wealth portfolio and the aggregate stock, we need to solve for the W/C as well as the P/D ratios as functions of the underlying state s_t . However, it would be technically implausible to obtain an analytical solution given there are two state variables (one exogenous X_t and one endogenous α_t) and the learning state evolution $d\alpha_t$ is non-trivial. In future empirical extension to this theory paper, we would resort to some numerical methods to jointly solve for the aggregate W/C ratio $\Lambda(X_t, \alpha_t)$ and the aggregate P/D ratio $\Phi(X_t, \alpha_t)$, taking into account some local pricing restrictions described in Appendix A7. Then, we could in turn illustrate the equity premium and the risk-free rate as functions of the state $s_t = (X_t, \alpha_t)$ numerically. We emphasize that while the agents are able to maintain a smooth consumption path, their W/C ratio jumps over time following changes in the underlying state of the world s_t , a property that also applies to the aggregate P/D ratio, capturing the fact that we do observe smooth sample paths of aggregate consumption even in the presence of stock market crashes (or some other forms of economic disasters/rare events that induce jump risk exposure) in reality.

⁹From the above analysis, aggregate stock return can be obtained as $R_t^S(s_t) = r_t^f(s_t) + EP_t^S(s_t)$ and the associated return dynamics $dR_t^S \equiv \frac{dP_t^S + D_t dt}{P_t^S}$ can be easily explored. Similarly, we can easily express the return dynamics of the aggregate wealth portfolio $dR_t^W \equiv \frac{dW_t + C_t dt}{W_t}$ as a function of the underlying state s_t .

4.5 Predictability Concern

In this section, we clarify some issues related to the consumption/dividend growth predictability as it was extensively debated over in the existing literature. Notably, one of the major criticisms against the long-run risk models points to their heavy reliance on a highly persistent aggregate consumption growth process, suggesting strong predictability of the aggregate consumption growth by valuation ratios, which is not true in the data. Typically, the standard long-run risk cookbook specifies the state evolution and consumption growth as $x_{t+1} = \rho x_t + \phi_\varepsilon \sigma \varepsilon_{t+1}$ and $g_{t+1} = \mu + x_t + \sigma \eta_{t+1}$ respectively, where the persistence parameter ρ plays a key role in achieving the major long-run pricing implications in this approach. Although our setup still has the flavor of a long-run model in which it still relies on recursive utility and non-i.i.d. consumption growth to pursue non-degenerate long horizon pricing implications, different from standard long-run risk models, valuation ratios in our model do not necessarily predict the aggregate consumption/dividend growth. We tackle this predictability issue by introducing an extra learning channel that produces endogenous uncertainty persistence, which in turn induces jumps in the valuation ratios, while keeping a level consumption with continuous sample paths, with no requirement for any persistent component coming from the state evolution, i.e. no imposed restrictions on strong predictability required for long horizon pricing. With this difference in mind, we point out that the long-run consumptions in our model are thus not risky. In the context of an ICAPM interpretation, we can regard the long-run dynamics as there are a lot of movements in the wealth portfolio but they correct themselves, i.e. things are safe in the long run. It is the long swings in the aggregate dividend that are important for the long-term investors and the current consumption becomes volatile as it responds to the short-term shifts in the wealth-consumption ratio induced by jumps in the underlying state of the world.

5 PRICING WITH LEARNING: MAIN THEORETICAL RESULTS

5.1 Dual Channels and Timing of Learning

In this model, besides the aggregate uncertainty stemming from jump risk exposure generated by structural changes in the aggregate consumption growth, the presence of Bayesian learning about the underlying economic state transitions also generates individual ambiguity coming from imperfect learning by the agents, which is the main novelty of our approach. To restate, here are the two major channels through which the equilibrium local risk prices will be affected:

1. channel of structural change in aggregate consumption growth induced by the jump state dependence of its drift and diffusion components;
2. channel of Bayesian learning by agents about the underlying state of the world at the next instant in time given all up-to-date information.

Below, we will illustrate in detail such dual-channel mechanics of the model, but before we do that, some clarification on the timing of the learning process shall be made. In general, agents in this economy perform Bayesian learning between any point in time t and the next instant t^+ . Given any state realization s at time t , agents learn and form a posterior update of the economic state transition intensity a_t^f based on all past information up to time t and any information updates between time t and time t^+ , before the realization of the time t^+ state s_{t^+} . After agents have learned and formed the posterior update a_t^f and accordingly adjusted their optimizing behavior, the time t^+ state s_{t^+} realizes and the time t^+ equity premium and risk-free rate are thus revealed. Such process continues through as agents learn continuously under constant uncertainty and thus the pricing moments get continuously revealed post-learning at each point in time.

5.2 Model Mechanism

To simplify illustration, recall that we are only looking at the case of simple learning with Markov structure on the bad state arrival as laid out in Section 3, where agents only learn about the bad state arrival intensity a_t while knowing exactly the bad state recovery intensity b . Also, without much loss of generality, in the following we will only look at the pricing implications for moments conditional on a good state realization at time t :¹⁰

$$EP_t^S(x_1, \alpha_t) = h\gamma x_1' \sigma_c \sigma_c' x_1 - a_t^f \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t^+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_2, \alpha_{t^+})}{\Phi(x_1, \alpha_t)} \right] \quad (25)$$

$$\begin{aligned} r_t^f(x_1, \alpha_t) &= \delta + \rho \left(x_1' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 \\ &\quad + \left(a_t^f - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t^+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} \right) - \frac{\theta - 1}{\theta} \left(a_t^f - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t^+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta} \right) \end{aligned} \quad (26)$$

where Λ is the W/C ratio of the representative agent and Φ is the P/D ratio of the aggregate stock. To sign the above conditional moments, we discuss several possible cases of preference parameter restrictions in the following.

Assumption 1. $\gamma > 1$ & $\rho < 1$ (**R1**) When $\rho = \frac{1}{EIS} < 1$, $EIS > 1$, the substitution effect dominates the wealth effect, and a reduction in investment opportunities (from good state to bad state in this case) in the form of lower expected consumption/dividend growth and higher volatility leads agents to reduce current consumption less than the reduction in investment. Hence, we should

¹⁰Intuitively, this can be justified by the fact that if we apply this to an economic disaster/rare event circumstance, the unconditional bad state arrival probability would be empirically so low (such events so rare) that the conditioning on a bad state realization can be suppressed without much loss of generality for the key pricing implications from the dual-channel mechanism. Similar arguments can be found in Gabaix (QJE 2012). Mechanically, this is straightforward in that if we condition on a bad state realization at time t , the time t learning state dynamics disappear and the pricing implications from the learning channel are degenerate.

expect locally $\Lambda(x_1, \cdot) > \Lambda(x_2, \cdot)$ and $\Phi(x_1, \cdot) > \Phi(x_2, \cdot)$, i.e. the W/C ratio conditional on current realization of a good state is higher than that conditional on current realization of a bad state, across all conditional belief updates (α_t) about the bad state arrival intensity. The same holds true for the P/D ratio since the aggregate stock is just a claim on levered consumptions (when the leverage $h = 1$, the two ratios coincide). Given these restrictions, we have $\theta < 0$, $\theta - 1 < 0$, and $1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} < 0$, $1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} > 0$, so

$$-a_t^f \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \right] > 0$$

i.e. the second term in $EP_t^S(x_1, \alpha_t)$ indeed captures an extra positive local pricing for the jump risk. Next, we can rewrite the last two terms of the local risk-free rate as

$$-a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} + \left(1 - \frac{1}{\theta}\right) a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta} + \frac{1}{\theta} a_t^f \equiv r_{ps,t} < 0$$

i.e. it indeed captures an extra local precautionary savings effect due to structural changes/regime switches in the underlying economic state of the world and thus a lower risk-free rate can be achieved.

Alternative Assumption. $\gamma > 1$ & $\rho > \gamma$ (**R2**) When $\rho = \frac{1}{EIS} > \gamma > 1$, $EIS < 1$, the wealth effect dominates the substitution effect, and a reduction in investment opportunities (from good state to bad state in this case) in the form of lower expected consumption/dividend growth and higher volatility leads agents to reduce current consumption more than the reduction in investment. Hence, we should expect locally $\Lambda(x_1, \cdot) < \Lambda(x_2, \cdot)$ and $\Phi(x_1, \cdot) < \Phi(x_2, \cdot)$, i.e. the W/C ratio conditional on current realization of a good state is lower than that conditional on current realization of a bad state, across all conditional belief updates (α_t) about the bad state arrival intensity. The same holds true for the P/D ratio. Thus, we have $\theta - 1 < 0$, $1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} < 0$ and $1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} > 0$. As a result, the second term in $EP_t^S(x_1, \alpha_t)$ still captures an extra positive local pricing for the jump risk. Again, rewrite the last two terms of the local risk-free rate as $r_{ps,t}$ and we find that $r_{ps,t} < 0$,¹¹ i.e. it again captures an extra local precautionary savings effect due to structural changes in the economy.

Alternative Assumption. $\gamma > 1$ & $1 < \rho < \gamma$ (**R3**) Again, $\rho = \frac{1}{EIS} > 1$, $EIS < 1$, the wealth effect dominates the substitution effect and a reduction in investment opportunities like before leads agents to reduce current consumption more than the reduction in investment. Hence, we should again expect locally $\Lambda(x_1, \cdot) < \Lambda(x_2, \cdot)$ and $\Phi(x_1, \cdot) < \Phi(x_2, \cdot)$, and thus $\theta - 1 > 0$,

¹¹ $\theta = \frac{1-\gamma}{1-\rho} > 0$ and $1 - \frac{1}{\theta} < 0$ imply $-a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} < 0$, $(1 - \frac{1}{\theta}) a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta} < 0$ and $-\frac{1}{\theta} (1 - a_t^f) < 0$.

$1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} < 0$, $1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} < 0$, and $-a_t^f \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1}\right] \left[1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)}\right] < 0$, i.e. the second term in $EP_t^S(x_1, \alpha_t)$ now captures a negative local pricing for the jump risk! If we look at the last two terms of the local risk-free rate, $r_{ps,t} \stackrel{\leq}{\geq} 0$, i.e. the extra precautionary savings effect cannot be established.¹²

Alternative Assumption. $\gamma = \rho$ (**R4, Power utility**) In this case, preferences degenerate to power utility. Although there is still inter-temporal risk from jump exposures to $\frac{\Lambda(x_2)}{\Lambda(x_1)}$ and $\frac{\Phi(x_2)}{\Phi(x_1)}$, the equilibrium pricing is independent of jumps because now agents do not have preferences over inter-temporal jump risk components:

$$\begin{aligned} EP_t^S(x_1) &= \gamma \sigma_c(x_1) \cdot \sigma_D(x_1) \\ &= \gamma \sigma_{cD}(x_1) \\ r_t^f(x_1) &= \delta + \gamma \left(x_1' \beta_c + \mu_c\right) - \frac{1}{2} \gamma (1 + \gamma) x_1' \sigma_c \sigma_c' x_1 \\ &= \delta + \gamma \mathbb{E}_t \frac{d \ln C_t}{dt} - \frac{1}{2} \gamma (1 + \gamma) \text{Var}_t(d \ln C_t) / dt \end{aligned}$$

This is exactly what we would have obtained under a standard C-CAPM model. Therefore, we emphasize that in order to generate meaningful equilibrium pricing implications in the context of a long-run model with event risk from a structural change/regime switch perspective, it is essential to have recursive utility, where we can separate inter-temporal tastes from risk aversion so that the agents can take both into account and fully adjust their optimal saving behavior.

Remark 4. *In what follows, we maintain (R1) as the benchmark assumption since the empirical literature suggest a larger than unity EIS. Thus, compared to Veronesi (JF 2000), even with $\gamma > 1$, higher uncertainty does increase risk premia, given a separated EIS from the risk aversion parameter.*

5.3 Perfect Information: Single Channel

First, let us ask the question what if there were no need to learn, i.e. agents were able to observe state transition intensities perfectly at true rates a and b ? In this case, the learning channel (α_t) is shut down and we are left with only the structural change/regime switch channel (X_t) of the aggregate consumption growth to affect equilibrium pricing. Intuitively, we now expect the model would collapse to one with implications that resemble those of Barro (2006), where the expected disaster/bad state arrival probability is constant over time. To see this, we write out the equity

¹² $\theta = \frac{1-\gamma}{1-\rho} > 0$ and $1 - \frac{1}{\theta} > 0$ imply $-a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta-1} < 0$, $(1 - \frac{1}{\theta}) a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}\right)^{\theta} > 0$ and $-\frac{1}{\theta} (1 - a_t^f) < 0$, so it is uncertain whether it is negative or not.

premium and the risk-free rate as follows, where the dependence on the extra learning state α_t previously seen is now gone.

$$EP_t^S(x_1) = h\gamma x_1' \sigma_c \sigma_c' x_1 - a \left[1 - \left(\frac{\Lambda(x_2)}{\Lambda(x_1)} \right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_2)}{\Phi(x_1)} \right] \quad (27)$$

$$\begin{aligned} r_t^f(x_1) &= \delta + \rho \left(x_1' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 \\ &\quad + \left(a - a \left(\frac{\Lambda(x_2)}{\Lambda(x_1)} \right)^{\theta-1} \right) - \frac{\theta - 1}{\theta} \left(a - a \left(\frac{\Lambda(x_2)}{\Lambda(x_1)} \right)^{\theta} \right) \end{aligned} \quad (28)$$

At a first glance, we notice that compared to traditional disaster models like Barro (2006), the aggregate consumption risk (induced by disaster/bad state arrivals) in this model does not manifest itself as direct jumps in the level of aggregate consumption, but rather through jumps in the W/C ratio induced by structural changes in the economic state/regime switches. In other words, jumps in Λ and thus $\frac{\Lambda(x_2)}{\Lambda(x_1)}$ (as well as $\frac{\Phi(x_2)}{\Phi(x_1)}$) are what matter now. As a result, even with jump risk to consumption, agents in this economy still enjoy a continuous sample path of level consumption. Notice that for the equity premium, the first term still captures the usual diffusion risk pricing, while the second term now captures the pricing for the jump risk induced by structural change/regime switch in the economy manifested as jumps in the W/C ratio as well as the P/D ratio. A closer look at the second term reveals a covariance-like relation between $\frac{\Lambda(x_2)}{\Lambda(x_1)}$ and $\frac{\Phi(x_2)}{\Phi(x_1)}$, where both ratios jump upon a bad state hits, with respective jump exposures captured by the two terms in the brackets. Intuitively, upon a bad state arrival at the next instant, value of the aggregate stock falls (jumps down) against its current dividend,¹³ which is captured by the positive jump risk exposure $1 - \frac{\Phi(x_2)}{\Phi(x_1)}$. Together with the negative term $1 - \left(\frac{\Lambda(x_2)}{\Lambda(x_1)} \right)^{\theta-1}$ coming from jumps in the SDF, which serves as the risk price device through the correlation with $\frac{\Phi(x_2)}{\Phi(x_1)}$, and the scaling by the constant transition intensity to a bad state, we confirm a positive risk premium contribution in the form of $-a \left[1 - \left(\frac{\Lambda(x_2)}{\Lambda(x_1)} \right)^{\theta-1} \right] \left[1 - \frac{\Phi(x_2)}{\Phi(x_1)} \right]$ that comes from the jump risk exposure induced by regime switches in this economy. This is intuitive because when a disaster/bad state hits, the aggregate stock market crashes, and agents (under *R1*) reduce holding equity more than their reduction in consumption, and therefore, to make the agents willing to still hold the aggregate equity, we have to compensate them for this jump risk by giving an extra premium on the aggregate stock return. In addition, as the structural consumption jump risk impacts on the inter-temporal savings behavior of the agents, the last two terms in equation (28) capture an extra precautionary savings effect (negative), and thus the local risk-free rate is now lower.

Proposition 1. *Under (R1) and Perfect Information, there is only one channel, i.e. the structural change/regime switch channel, which induces higher (conditional) equity premium and lower (condi-*

¹³Dividends of the aggregate stock also have a continuous sample path like the aggregate consumption.

tional) risk-free rate in equilibrium, in the form of positive compensation for jump risk exposure and negative precautionary savings effect.

5.4 Imperfect Information: Dual Channels

Now, let us switch on the learning channel by noting the fact that agents in this economy do not fully observe the true transition intensity a of a bad state arrival and thus have to form and continuously update their beliefs a_t^f through imperfect learning over time. In this case, the conditional equity premium and risk-free rate take the form as in equations (25) and (26). Learning enters pricing through the dependence of the W/C and P/D ratios on not only the economic state that causes structural changes in aggregate consumption, but also the learning state that updates beliefs about conditional bad state arrival given imperfect information. Thus, to understand the pricing impact of learning, we need to examine the properties of Λ and Φ as functions of both X_t and α_t .

Assumption 2. *Monotonicity and Convexity (M&C)*

$$\frac{\partial \Lambda(x, \alpha)}{\partial \alpha} > 0 \text{ and } \frac{\partial^2 \Lambda(x, \alpha)}{\partial \alpha^2} > 0, \forall x$$

i.e. the W/C ratio is monotonic increasing and strictly convex in α , across all economic state realizations x . And, given the aggregate dividend follows a levered consumption process, $\Phi(x, \alpha)$ should have the same property.

Intuitively, conditional on being in any given state x_1 (good) or x_2 (bad), when the agents believe it is more likely for them to face a low bad state arrival intensity (a_l), i.e. $\alpha \uparrow$, conditioning on all up-to-date information, it suggests a more optimistic view towards future economic performance, and the agents would expand their investments and thus push up the value of aggregate wealth (a claim on future aggregate consumptions) against current consumption since it is now more likely to go into a good time at the next instant when expected consumption growth is higher and volatility is lower. And vice versa for a decrease in α which suggests a pessimistic view prevails. More importantly, such marginal effect is increasing in α , i.e. the W/C ratio is convex in agents' beliefs. This is where the learning-induced individual ambiguity aversion kicks in, and the intuition is as follows.

5.4.1 Aggregate Uncertainty and Individual Ambiguity

When α is close to 1 or 0, it is almost certain to the agents that the conditional bad state arrival intensity is a_l or a_h , and therefore, although it is still uncertain whether a bad or a good economic state will arrive at the next instant (still aggregate uncertainty), the agents are much less ambiguous about the odds of getting into either states (less individual ambiguity). On the other hand, when

α is close to 0.5, agents have almost no clue about how likely it is for them to face a low or a high conditional bad state arrival intensity, and thus, the ambiguity to the agents towards the odds of getting into either economic states is approaching the greatest. Axiomatically, we have assumed that agents in this world are not only risk averse but also ambiguity averse. Then, the intuition of the convexity assumption naturally follows. Say α_t was initially close to 0 (agents were almost certain that a bad state is going to come with high probability), now if α_t rises, this implies a lower posterior belief update about the expected conditional bad state arrival intensity a_t^f , conditioning on time t information history. Yet, at the same time, as a higher α_t now becomes closer to 0.5, which raises the agents' ambiguity about the odds (a_t^f) of getting into either economic states, the two channels work in opposite directions, generating opposite views towards future economic performance, which in turn makes the agents expand investments less and thus push up the value of aggregate wealth against current consumption less than in a model without ambiguity and more generally without learning. On the other hand, say α_t was initially close to 1 (agents were almost certain that a bad state is going to come with low probability), now if α_t falls, this implies a higher posterior belief update about the expected conditional bad state arrival intensity a_t^f . But at the same time, as a lower α_t now becomes closer to 0.5, the agents are more ambiguous about the odds (a_t^f) of getting into either economic states. Thus, in this case, both channels work in the same direction to generate a more pessimistic view towards future economic performance, which in turn incentivizes the agents to reduce investments and thus trade down the value of aggregate wealth against current consumption more than in a model without learning. Combining the above, and given that without learning, a time-varying disaster rate (result of some other causes) produces a linear (or quasi-linear) relation between the W/C ratio and the disaster rate itself (e.g. Gabaix, QJE 2012), the learning behavior and associated ambiguity aversion in this model enriches the dynamics by making such a relation convex.¹⁴

5.4.2 Dual Puzzles and Regime Learning

To have a closer look at learning's impact on equilibrium pricing, we decompose the ratio of Λ 's into a product of two terms:

$$\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} = \frac{\Lambda(x_2, \alpha_t)}{\Lambda(x_1, \alpha_t)} \frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_2, \alpha_t)}$$

where the first term corresponds to the case without learning, while the second term captures the impact of learning facing the arrival of a bad economic state. From previous discussion, under Assumption 1, $\frac{\Lambda(x_2, \alpha_t)}{\Lambda(x_1, \alpha_t)} < 1$, and recall from Section 3 on the learning state dynamics ($d\alpha_t$), facing a bad state arrival at the next instant, $\alpha_{t+} < \alpha_t$, so $\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_2, \alpha_t)} < 1$. As a result, with learning behavior, the ratio $\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}$ is even smaller, making the last two terms in equation (26) even more negative.

¹⁴A similar mechanism was first proposed by Veronesi (JF 2000) in a model where agents learn about the drift of the aggregate dividend process.

Therefore, we obtain a stronger effect (than standard event risk models without learning) from the precautionary savings motive on the local risk-free rate so that it can be suppressed even lower in equilibrium. Hence, learning helps improve the resolution of low risk-free rate observed. As for the equity premium, equation (25), we can do a similar decomposition:

$$\frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} = \frac{\Phi(x_2, \alpha_t)}{\Phi(x_1, \alpha_t)} \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_2, \alpha_t)}$$

where the first term corresponds to the case without learning, while the second term captures the impact of learning facing the arrival of a bad state. Since the aggregate dividend follows a levered consumption process, properties of Λ apply one-for-one to Φ . So, we should also have $\frac{\Phi(x_2, \alpha_t)}{\Phi(x_1, \alpha_t)} < 1$ and $\frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_2, \alpha_t)} < 1$, which together give us an even higher $1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)}$ than in the absence of learning. Moreover, as we just showed that $\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)}$ is now smaller with learning, $- \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} \right]$ is even larger than before. Put together, these results lead to a higher equilibrium (conditional) equity risk premium compensated for the jump risk exposure in the presence of learning. Thus, learning works towards improving explanations for a high equity premium observed.

Proposition 2. *Under (R1) and Imperfect Information, both the structural change/regime switch channel and the learning channel work to induce higher (conditional) equity premium and lower (conditional) risk-free rate in equilibrium, in the form of positive compensation for jump risk exposure and negative precautionary savings effect (of larger magnitudes in the presence of learning).*

6 CONCLUSION

In this paper, rather than treating event risk as direct jumps in the aggregate consumption process, we model it as changes in the underlying state of the world, which affects consumption and dividend flows through their growth and volatility dependence on the state. Information about the state transition is imperfect in this economy and agents perform Bayesian learning to form and update beliefs about the conditional state arrival to guide them towards dynamic optimization. As we saw, this new learning component in the time-varying jump risk modeling approach generates novel pricing implications through inducing an extra covariance to be priced in equilibrium. Specifically, besides the aggregate uncertainty stemming from jump risk exposure generated by structural changes in the aggregate consumption process, the presence of Bayesian learning behavior generates individual ambiguity coming from imperfect learning by the agents. As we have analyzed above, such dual channels can help better explain theoretically the dual puzzles of asset prices, resolve major predictability issues, and address the time-varying conditional first and second pricing moments. Overall, this is a theoretical modeling paper. In future empirical research, we would try to calibrate

the model herein to obtain consumption and learning parameters, after which we shall numerically solve for Λ and Φ as functions of x and α .¹⁵ With a fully calibrated model, we can then go on to try matching some important moments of observed asset prices to evaluate the empirical relevance of this modeling approach. At the same time, we can further compare the empirical performance of this imperfect learning-induced time-varying jump risk modeling framework for consumption-based asset pricing with some other established models along the line such as the old disaster models, e.g. Rietz (1988) and Barro (2006), the long-run models, e.g. Bansal-Yaron (2004) and Campbell-Vuolteenaho (2004), the habit models, e.g. Campbell-Cochrane (1999) and Santos-Veronesi (2009), and the new disaster models, e.g. Gabaix (2012) and Wachter (2012). In addition, we can potentially extend this framework to generalized learning, i.e. learning about both good state and bad state transition intensities, both arrivals and recoveries, or in an n -state model, learning about the transitions across all states (learning about the entire transition matrix). Last but not least, we might be able to extend this modeling framework to alternative specifications of the learning process, e.g. a filtering problem of bad state arrivals from the past history, to alternative specifications of the aggregate dividend process, e.g. an independent diffusion process, or to alternative specifications of the information structure that are subject to learning, e.g. learning about return predictability on top of transition intensity learning, all of which might be used as a checking device for the robustness of the learning-induced results obtained herein.

¹⁵We can then graph the simulated Λ against α for fixed values of x , to confirm the monotonicity and convexity assumption, and also confirm that for fixed α and with appropriate boundary conditions considered, locally $\Lambda(x_1, \cdot) > \Lambda(x_2, \cdot)$. The same goes for Φ . Since aggregate dividend is a levered consumption claim, Φ on average should be more curved than Λ . We can also illustrate the dual channels by graphing Λ over time, as α continuously updates and x jumps. The same goes for Φ .

APPENDIX: DERIVATIONS AND PROOFS

$$\begin{aligned} \mathbf{A1.} \quad d\alpha_t/dt &= \xi_t + \kappa_t \alpha_t = \bar{\alpha} \left(a_t^f - a_l \right) + \frac{a_l - a_t^f}{a_t^f} \bar{\alpha} = \bar{\alpha} \left(\frac{a_t^f a_t^f - a_l a_t^f}{a_t^f} + \frac{a_l - a_t^f}{a_t^f} \right) \\ &= \frac{\bar{\alpha}}{a_t^f} \left[a_t^f \left(a_t^f - 1 \right) - a_l \left(a_t^f - 1 \right) \right] = \frac{\bar{\alpha}}{a_t^f} \left(a_t^f - a_l \right) \left(a_t^f - 1 \right) < 0 \end{aligned}$$

$$\mathbf{A2.} \quad \frac{W_t}{C_t} \equiv \frac{\frac{\partial V_t}{\partial C_t}}{C_t} = \frac{V_t^{1-\rho}}{(1-\exp(-\delta\varepsilon))C_t^{1-\rho}} = (1 - \exp(-\delta\varepsilon))^{-1} \left(\frac{V_t}{C_t} \right)^{1-\rho}$$

A3. Eigenvalue Problem and Eigenfunction ψ

In fact, we can think of equation (13) as an eigenvalue problem with positive eigenfunction ψ and related eigenvalue $\eta = \frac{\delta(1-\gamma)}{1-\rho}$:

$$\lim_{\varepsilon \downarrow 0} \frac{\mathbb{E} \left[\psi \left(s_{t+\varepsilon} \right) \left(\frac{C_{t+\varepsilon}}{C_t} \right)^{1-\gamma} \middle| s_t = s \right] - \psi(s)}{\varepsilon} = \frac{\delta(1-\gamma)}{1-\rho} \psi(s) = \eta \psi(s)$$

Note that if η is an eigenvalue to the functional equation above, with related eigenfunction ψ that is positive on its domain, then the condition $\delta > \frac{\eta(1-\rho)}{1-\gamma}$ guarantees the existence of a solution to the recursive utility formulation. Thus, to solve for the SDF dynamics, let $M_{t+\varepsilon} \equiv \left(\frac{C_{t+\varepsilon}}{C_t} \right)^{1-\gamma}$, $M_t = \left(\frac{C_t}{C_0} \right)^{1-\gamma}$, and $d \ln M_t = (1-\gamma) d \ln C_t = (1-\gamma)(X_t' \beta_c + \mu_c) dt + (1-\gamma) X_t' \sigma_c dW_t$. By defining the generator for a distorted conditional expectation process $\mathbb{B} \equiv \mathbb{E}_t[M_{t+} f_{t+}]$ and driving δ to its lower limit $\delta_{inf} = \frac{\eta(1-\rho)}{1-\gamma}$, we can manipulate to obtain the functional eigenvalue problem in the following matrix form

$$\mathbb{B} \psi(s) = \eta \psi(s) \quad \text{or} \quad \mathbb{E} [M_t \psi(s_t) | s_0 = s] = \exp(\eta t) \psi(s)$$

where η is an eigenvalue to the above problem. Now we guess the eigenfunction is log-linear in the state and we can solve for $\psi(s_t) = \psi(X_t, \alpha_t)$.

A4. Deriving $dlnV_t^*$ and the SDF dynamics

Suppose we guess the deterministic drift component of $dlnV_t^*$ is of the form $\beta_v^* \cdot X_t + \mu_v^*$, then we should have the following SDF process:

$$\begin{aligned}
 dlnS_t &= \left(-\delta - \rho X_t' \beta_c - \rho \mu_c \right) dt - \rho X_t' \sigma_c dW_t + \frac{\rho - \gamma}{1 - \gamma} [(\beta_v^* \cdot X_t + \mu_v^*) dt + \xi_v^*(X_t) dW_t + \chi_v^*(s_{t+}, s_t)] \\
 &= \left(-\delta - \rho X_t' \beta_c - \rho \mu_c \right) dt - \rho X_t' \sigma_c dW_t \\
 &\quad + \frac{\rho - \gamma}{1 - \gamma} \left[(\beta_v^* \cdot X_t + \mu_v^*) dt + (1 - \gamma) X_t' \sigma_c dW_t + \ln \frac{\psi(s_{t+})}{\psi(s_t)} \right] \\
 &= \underbrace{\left[-\delta - \rho X_t' \beta_c - \rho \mu_c + \frac{\rho - \gamma}{1 - \gamma} (\beta_v^* \cdot X_t + \mu_v^*) \right]}_{\equiv \beta_s(s_t)} dt + \underbrace{\left[-\rho X_t' \sigma_c + (\rho - \gamma) X_t' \sigma_c \right]}_{\equiv \xi_s(s_t) = -\gamma X_t' \sigma_c} dW_t \\
 &\quad + \underbrace{\frac{\rho - \gamma}{1 - \gamma} \ln \frac{\psi(s_{t+})}{\psi(s_t)}}_{\equiv \chi_s(s_{t+}, s_t)} \\
 dlnS_t &= \beta_s(s_t) dt + \xi_s(s_t) dW_t + \chi_s(s_{t+}, s_t)
 \end{aligned}$$

To complete the description of the SDF process, we must solve out the drift of $dlnV_t^*$ in terms of the underlying state s_t . To do this, first note that the generator \mathbb{B} defined before can be expressed in matrix form as $\mathbb{B} = A^f + D$, where $A^f \equiv \begin{bmatrix} -a^f & a^f \\ b & -b \end{bmatrix}$ and $D \equiv (1 - \gamma) (\text{diag}(\beta_c) + \mu_c I_2) + \frac{1}{2} (1 - \gamma)^2 \text{diag}(\sigma_c \sigma_c')$. Next, let $K \equiv [\exp(\chi_v^*(s_+, s))]_{2 \times 2}$, $D^* \equiv \text{diag}(\beta_v^*) + \mu_v^* I_2 + \frac{1}{2} (1 - \gamma)^2 \text{diag}(\sigma_c \sigma_c')$, and we can thus define a new generator $\mathbb{B}^* \equiv K \circ A^f + D^*$ under the equivalent martingale measure with respect to $dlnV_t^*$.¹⁶ Hence, by the martingale property of V_t^* , we must have $\mathbb{B}^* \mathbf{1}_2 = \mathbf{0}_2$, i.e. the local pricing restriction from risk-return trade-off.¹⁷ Thus, given a solution of $\psi(s_t)$, we can back out the expression for $\beta_v^* \cdot X_t + \mu_v^*$ through the following

$$\begin{aligned}
 \mathbf{0}_2 &= (K \circ A^f + D^*) \mathbf{1}_2 \\
 \mathbf{0}_2 &= \left\{ [\exp(\chi_v^*(s_{t+}, s_t))] \circ \begin{bmatrix} -a_t^f & a_t^f \\ b & -b \end{bmatrix} + \text{diag}(\beta_v^*) + \mu_v^* I_2 + \frac{1}{2} (1 - \gamma)^2 \text{diag}(\sigma_c \sigma_c') \right\} \mathbf{1}_2 \\
 \mathbf{0}_2 &= \begin{bmatrix} -a_t^f & a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \\ b \frac{\psi(x_1, \alpha_{t+})}{\psi(x_2, \alpha_t)} & -b \end{bmatrix} \mathbf{1}_2 + \text{diag}(\beta_v^*) \mathbf{1}_2 + \mu_v^* I_2 \mathbf{1}_2 + \frac{1}{2} (1 - \gamma)^2 \text{diag}(\sigma_c \sigma_c') \mathbf{1}_2
 \end{aligned}$$

¹⁶We have suppressed the time subscript here, “ \circ ” denotes the element-wise product, “ $\text{diag}(\sim)$ ” means creating a diagonal matrix with entries taken from \sim , and $s = (x, a^f)$.

¹⁷This is an application of the Girsanov’s Theorem.

which collapses down to

$$\begin{aligned} -a_t^f + a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} + x_1' \beta_v^* + \mu_v^* + \frac{1}{2} (1 - \gamma)^2 x_1' \sigma_c \sigma_c' x_1 &= 0 \\ -b + b \frac{\psi(x_1, \alpha_{t+})}{\psi(x_2, \alpha_t)} + x_2' \beta_v^* + \mu_v^* + \frac{1}{2} (1 - \gamma)^2 x_2' \sigma_c \sigma_c' x_2 &= 0 \end{aligned}$$

which defines $\beta_v^* \cdot X_t + \mu_v^*$ as a function of the underlying state of the world. Therefore, the SDF follows the process

$$\begin{aligned} d \ln S_t &= \beta_s(s_t) dt + \xi_s(s_t) dW_t + \chi_s(s_{t+}, s_t) \\ \text{where } \beta_s &= -\delta - \rho X_t' \beta_c - \rho \mu_c + \frac{\rho - \gamma}{1 - \gamma} (X_t' \beta_v^* + \mu_v^*) \\ \xi_s &= -\gamma X_t' \sigma_c \\ \chi_s(s_{t+}, s_t) &= \ln \left(\frac{\psi(s_{t+})}{\psi(s_t)} \right)^{\frac{\theta-1}{\theta}} \\ X_t' \beta_v^* + \mu_v^* &= \begin{cases} x_1' \beta_v^* + \mu_v^* = a_t^f - a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} - \frac{1}{2} (1 - \gamma)^2 x_1' \sigma_c \sigma_c' x_1, & X_t = x_1 \\ x_2' \beta_v^* + \mu_v^* = b - b \frac{\psi(x_1, \alpha_{t+})}{\psi(x_2, \alpha_t)} - \frac{1}{2} (1 - \gamma)^2 x_2' \sigma_c \sigma_c' x_2, & X_t = x_2 \end{cases} \end{aligned}$$

A5. Deriving local risk premia

$$\begin{aligned} r_t^e(x_1, \alpha_t) &= -\xi_s(s_t) \cdot \xi_{PS}(s_t) - \left(\begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} -a_t^f & a_t^f \\ b & -b \end{bmatrix} \right) \\ &\quad \left[\begin{array}{c} \begin{bmatrix} 1 - \left(\frac{\psi(x_1, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \\ 1 - \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \end{bmatrix} \begin{bmatrix} 1 - \frac{\Phi(x_1, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \\ 1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \end{bmatrix} \\ \begin{bmatrix} 0 \\ 1 - \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \end{bmatrix} \begin{bmatrix} 1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \end{bmatrix} \end{array} \right] \\ &= \gamma x_1' \sigma_c \cdot \sigma_D(x_1) - \begin{bmatrix} -a_t^f & a_t^f \end{bmatrix} \left[\begin{array}{c} 0 \\ 1 - \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \end{array} \right] \begin{bmatrix} 1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \end{bmatrix} \\ EP_t^S(x_1, \alpha_t) &= \gamma x_1' \sigma_c \cdot h x_1' \sigma_c - a_t^f \left[1 - \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \right] \begin{bmatrix} 1 - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \end{bmatrix} \end{aligned}$$

$$\text{and similarly } EP_t^S(x_2, \alpha_t) = \gamma x_2' \sigma_c \cdot h x_2' \sigma_c - b \left[1 - \left(\frac{\psi(x_1, \alpha_t)}{\psi(x_2, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \right] \begin{bmatrix} 1 - \frac{\Phi(x_1, \alpha_t)}{\Phi(x_2, \alpha_t)} \end{bmatrix}.$$

A6. Deriving local risk-free rates:

Local risk-free rates can be obtained from the parametrization of the SDF as functions of the underlying state $s_t = (X_t, \alpha_t)$ as follows.

$$\begin{aligned} r_t^f(X_t, \alpha_t) &= -\hat{\mathbb{B}}1_2 \\ &= -\left(\hat{K} \circ A_t^f + \hat{D}\right)1_2 \\ &= -\left[\exp(\chi_s(s_{t+}, s_t))\right] \circ \begin{bmatrix} -a_t^f & a_t^f \\ b & -b \end{bmatrix} + \text{diag}(\beta_s) + \frac{1}{2}\text{diag}(\xi_s \xi_s') \Big] 1_2 \end{aligned}$$

and

$$\begin{aligned} x_1' \beta_v^* + \mu_v^* &= a_t^f - a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} - \frac{1}{2}(1-\gamma)^2 x_1' \sigma_c \sigma_c' x_1 \\ x_2' \beta_v^* + \mu_v^* &= b - b \frac{\psi(x_1, \alpha_t)}{\psi(x_2, \alpha_t)} - \frac{1}{2}(1-\gamma)^2 x_2' \sigma_c \sigma_c' x_2 \end{aligned}$$

When the economy is in the good state, i.e. $X_t = x_1$, we have

$$\begin{aligned} \beta_s &= -\delta - \rho x_1' \beta_c - \rho \mu_c + \frac{\rho - \gamma}{1 - \gamma} (x_1' \beta_v^* + \mu_v^*) \\ \xi_s \xi_s' &= \gamma^2 x_1' \sigma_c \sigma_c' x_1 \\ \chi_s(s_{t+}, s_t) &= \ln \left(\frac{\psi(s_{t+})}{\psi(s_t)} \right)^{\frac{\theta-1}{\theta}} \end{aligned}$$

and thus we can plug in such parameters to obtain

$$\begin{aligned}
r_t^f(x_1, \alpha_t) &= - \begin{bmatrix} 1 \\ 0 \end{bmatrix}' \left[\begin{bmatrix} \exp(\chi_s(s_1, s_1)) & \exp(\chi_s(s_2, s_1)) \\ \exp(\chi_s(s_1, s_2)) & \exp(\chi_s(s_2, s_2)) \end{bmatrix} \circ A_t^f + \text{diag}(\beta_s) + \frac{1}{2} \text{diag}(\xi_s \xi_s') \right] 1_2 \\
&= - \begin{bmatrix} 1 \\ 0 \end{bmatrix}' \begin{bmatrix} -a_t^f & a_t^f \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} \\ b \left(\frac{\psi(x_1, \alpha_t)}{\psi(x_2, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} & -b \end{bmatrix} 1_2 \\
&\quad - \begin{bmatrix} 1 \\ 0 \end{bmatrix}' \text{diag}(\beta_s) 1_2 - \frac{1}{2} \begin{bmatrix} 1 \\ 0 \end{bmatrix}' \text{diag}(\xi_s \xi_s') 1_2 \\
&= a_t^f - a_t^f \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} + \delta + \rho x_1' \beta_c + \rho \mu_c - \frac{\rho - \gamma}{1 - \gamma} (x_1' \beta_v^* + \mu_v^*) - \frac{1}{2} \gamma^2 x_1' \sigma_c \sigma_c' x_1 \\
&= \delta + \rho (x_1' \beta_c + \mu_c) + \frac{1}{2} (\rho - \gamma - \rho \gamma) x_1' \sigma_c \sigma_c' x_1 \\
&\quad + a_t^f - a_t^f \left(\frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right)^{\frac{\theta-1}{\theta}} - \frac{\theta - 1}{\theta} \left(a_t^f - a_t^f \frac{\psi(x_2, \alpha_{t+})}{\psi(x_1, \alpha_t)} \right) \\
&= \delta + \rho (x_1' \beta_c + \mu_c) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 + a_t^f \\
&\quad - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} - \frac{\theta - 1}{\theta} \left(a_t^f - a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^\theta \right)
\end{aligned}$$

and similarly when the economy is in the bad state, i.e. $X_t = x_2$, we have

$$r_t^f(x_2, \alpha_t) = \delta + \rho (x_2' \beta_c + \mu_c) - \frac{1}{2} \gamma (1 + \rho) x_2' \sigma_c \sigma_c' x_2 + b - b \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^{\theta-1} - \frac{\theta-1}{\theta} \left(b - b \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^\theta \right).$$

A7. Local Restrictions on Λ and Φ

From equations (11) and (12), we can obtain

$$\begin{aligned}
0 &= -\frac{\delta}{1 - \rho} [(1 - \exp(-\delta\varepsilon)) \Lambda(s_t)] + \frac{1}{1 - \gamma} [(1 - \exp(-\delta\varepsilon)) \Lambda(s_t)]^{1-\theta} \\
&\quad \cdot \lim_{\varepsilon \downarrow 0} \frac{\mathbb{E}_t \left[(1 - \exp(-\delta\varepsilon))^\theta \Lambda(s_{t+\varepsilon})^\theta M_{t+\varepsilon} \right] - (1 - \exp(-\delta\varepsilon))^\theta \Lambda(s_t)^\theta}{\varepsilon}
\end{aligned}$$

Restriction 1. *The above (differential) equation can be manipulated (by Ito's lemma with jumps) into the following local restrictions on the aggregate wealth-consumption ratio $\Lambda(X_t, \alpha_t)$ and the*

aggregate price-dividend ratio $\Phi(X_t, \alpha_t)$.

$$\begin{aligned} \frac{C_t}{W_t} \Big|_{x_1} = \Lambda(x_1, \alpha_t)^{-1} &= \delta + \rho \left(x_1' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 \\ &\quad - \xi_t \frac{\partial \ln \Lambda(x_1, \alpha_t)}{\partial \alpha} + a_t^f \frac{1}{\theta} \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^\theta \right] \end{aligned} \quad (29)$$

$$\begin{aligned} \frac{C_t}{W_t} \Big|_{x_2} = \Lambda(x_2, \alpha_t)^{-1} &= \delta + \rho \left(x_2' \beta_c + \mu_c \right) - \frac{1}{2} \gamma (1 + \rho) x_2' \sigma_c \sigma_c' x_2 \\ &\quad + b \frac{1}{\theta} \left[1 - \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^\theta \right] \end{aligned} \quad (30)$$

$$\begin{aligned} \frac{D_t}{P_t^S} \Big|_{x_1} = \Phi(x_1, \alpha_t)^{-1} &= \delta + \rho \left(x_1' \beta_c + \mu_c \right) - \mu_D(x_1) - \frac{1}{2} \gamma (1 + \rho) x_1' \sigma_c \sigma_c' x_1 + \gamma \sigma_c(x_1) \sigma_D(x_1) \\ &\quad - \xi_t \frac{\partial \ln \Phi(x_1, \alpha_t)}{\partial \alpha} + a_t^f \frac{1}{\theta} \left[1 - \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^\theta \right] \\ &\quad + a_t^f \left(\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} \right)^{\theta-1} \left[\frac{\Lambda(x_2, \alpha_{t+})}{\Lambda(x_1, \alpha_t)} - \frac{\Phi(x_2, \alpha_{t+})}{\Phi(x_1, \alpha_t)} \right] \end{aligned} \quad (31)$$

$$\begin{aligned} \frac{D_t}{P_t^S} \Big|_{x_2} = \Phi(x_2, \alpha_t)^{-1} &= \delta + \rho \left(x_2' \beta_c + \mu_c \right) - \mu_D(x_2) - \frac{1}{2} \gamma (1 + \rho) x_2' \sigma_c \sigma_c' x_2 + \gamma \sigma_c(x_2) \sigma_D(x_2) \\ &\quad + b \frac{1}{\theta} \left[1 - \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^\theta \right] \\ &\quad + b \left(\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} \right)^{\theta-1} \left[\frac{\Lambda(x_1, \alpha_t)}{\Lambda(x_2, \alpha_t)} - \frac{\Phi(x_1, \alpha_t)}{\Phi(x_2, \alpha_t)} \right] \end{aligned} \quad (32)$$

(Note that these restrictions will be used in the numerical solution methods for the two critical aggregate ratios in equilibrium.)

References

- [1] Barro, Robert, 2006. "Rare Disasters and Asset Markets in the Twentieth Century," Scholarly Articles 3208215, Harvard University Department of Economics.
- [2] Bansal, Ravi & Yaron, Amir, 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, American Finance Association, vol. 59(4), pages 1481-1509, 08.
- [3] Campbell, John Y. & Vuolteenaho, Tuomo, 2004. "Bad Beta, Good Beta," *American Economic Review*, American Economic Association, vol. 94(5), pages 1249-1275, December.
- [4] Epstein, Larry G. & Zin, Stanley E., 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, University of Chicago Press, vol. 99(2), pages 263-86, April.
- [5] Gabaix, Xavier, 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *The Quarterly Journal of Economics*, Oxford University Press, vol. 127(2), pages 645-700.
- [6] Hansen, Lars P. & Scheinkman, Jose A., 2012. "Recursive Utility in a Markov Environment with Stochastic Growth," Working Papers 2012-002, Becker Friedman Institute for Research In Economics.
- [7] Liptser, Robert S. & Shiryaev, Albert N., 2001. *Statistics of Random Processes*, translated by A. B. Aries; translation editor, Stephen S. Wilson. Edition: 2nd, rev. and expanded ed.. Imprint: Berlin; New York: Springer, c2001.
- [8] Rietz, Thomas A., 1988. "The equity risk premium a solution," *Journal of Monetary Economics*, Elsevier, vol. 22(1), pages 117-131, July.
- [9] Veronesi, Pietro, 2000. "How does Information Quality Affect Stock Returns?" *Journal of Finance*, 55(2), April 2000.
- [10] Wachter, Jessica A. , Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? (August 7, 2012). *Journal of Finance*, Forthcoming; AFA 2011 Denver Meetings Paper.

-END-