Internet Appendix to "Time-Varying Ambiguity, Credit Spreads, and the Levered Equity Premium"

Zhan Shi

Tsinghua PBCSF

November 16, 2018

This internet appendix includes four sections. Section IA.A presents a model of learning under ambiguity to justify the ambiguity dynamics as exogenously specified in Section 2.2. Specifically, Monte Carlo simulations based on the learning model demonstrate that the exogenous process in Eq. (10) captures several important properties of temporal variation in the ambiguity level: it follows a persistent but stationary process, and the conditional volatility of ambiguity shocks varies over time and positively comoves with the ambiguity level. The key to these properties is the prior-by-prior updating of the belief set under the multiple-priors utility. It leads to asymmetric treatments to ambiguous information such that the ambiguity does not vanish in the long run and holds many desired dynamics properties. Our simulations also derive important quantitative implications for the model-implied ambiguity process, which closely coincide with those carried by our empirical measure.

More details on this ambiguity measure are covered in Section IA.B. Section IA.B.1 explains why the dataset provided by Blue Chip Financial Forecasts (BCFF) is most suitable for constructing an empirical proxy for aggregate ambiguity and investigating its role in asset pricing. In Section IA.B.2, we show that our ambiguity measure exhibits time-varying volatility and a right-skewed distribution, which actually motivate the CIR-type specification in Eq. (10). These empirical properties are also consistent with the ambiguity processes simulated from our learning model. Therefore, the first two sections in this internet appendix jointly demonstrate that unity between the learning model, the exogenous ambiguity dynamics, and the data of survey forecasts can be reached. Finally, Section IA.B.3 investigates the robustness of $\widetilde{A}_t$'s predictive power when other return predictors (not examined in Section 3.2) are controlled for.

Results in Section IA.C build on a sample of firms that defaulted on their debt obligations. This section serves two purposes: estimating the cost of default and measuring the default boundary. The former is fundamental to our model calibration, in which we need to specify how the default cost co-varies with the ambiguity level, and the latter is required to verify the model's prediction that the (value-based) boundary level decreases with the degree of ambiguity. In this appendix, we describe the data sources and present the summary statistics of the estimated default boundary/cost; our estimates turn out highly consistent with those reported in Davydenko (2012) and Davydenko et al. (2012).

Finally, Section IA.D provides details on how to solve for equity and debt values by regular perturbation. And Section IA.E shows that the "medians-to-medians" approach adequately addresses the issue of convexity bias in credit spreads.

Appendix IA.A. A Model of Learning under Ambiguity

The purpose of this appendix is to show that a model in the spirit of Epstein and Schneider (2007) and Illeditsch (2011) can endogenously generate an ambiguity process that resembles the one specified exogenously in Section 2.2. This model extends the aforementioned theoretical works in terms of the definition of ambiguous information quality: the agent is uncertain not only about the precision of signals but also about the relevance of signals. Moreover, the model closely mimics the information processing in practice by an ambiguity-averse agent, who tries to learn the conditional consumption growth rate but does not have perfect and timely observations about the realized growth rate. As such, it builds a direct link between Knightian uncertainty perceived by the agent and the disagreement among experts; the latter is utilized in Section 3.1 for the inference of historical ambiguity levels.

IA.A.1. Model Setup

The true data-generating process for aggregate consumption is modeled as follows:[IA.1]

$$dc_t = (\bar{\mu} - \frac{1}{2}\sigma_c^2 + x_t)dt + \sigma_c dB_{c,t} \qquad \text{(IA.A1)}$$

$$dx_t = -\kappa_x x_t dt + \sigma_c dB_{x,t}. \qquad \text{(IA.A2)}$$

The time-varying expected growth rate $\bar{\mu} + x_t$ follows the specification of Campbell (1999, 2003). Note that it is not necessarily as persistent as the process typically calibrated in studies of long-run risk models. For example, Piazzesi and Schneider (2006) adopt a similar specification, and their estimates (based on GMM) of κ_x range from 0.396 to 0.456. In the following simulation exercise, we infer the dynamics of x_t from consensus survey forecasts and let the data reveal the degree of persistence.

As is standard in the literature of learning models, the expected growth rate is unobservable and needs to be filtered out using the realized growth rate. However, our model also takes into account the fact that investors are highly unlikely to learn the true level of consumption in real

[IA.1] The symbols defined as follows apply only to this appendix.

time. First of all, the NIPA consumption data pertaining to quarter i is regularly released during the second month of quarter $i+1$. Second, this release is termed a "preliminary release" and is typically based on projections; it is subject to revisions not only in subsequent months, as BEA statisticians have more and more complete sample data to perform imputations and interpolations, but also in a few years, e.g., because of changes in the procedures of seasonal adjustments. Finally, even if we view this filtering problem from the perspective of econometricians, who can observe a long time series of historical consumption expenditures from the BEA, ample evidence has been found for the imperfections in the measurement of NIPA consumption (see, e.g., Wilcox (1992)). Given these practical considerations, we model both x_t and c_t as latent state variables. The agent, on the other hand, receives noisy public information about c_t with a lag

$$dy_t = c_t dt + \sigma_y dB_{y,t}. \tag{IA.A3}$$

The term y_t can be interpreted as the preliminary or the second release of real consumption expenditures from the BEA. Besides the official releases of consumption growth, the agent also collects intangible information for the inference of c_t and x_t. This intangible information is modeled as private signals that are correlated with innovations in expected and realized consumption growth:

$$\begin{aligned} d\nu_t &= \sigma_\nu \phi_c dB_{c,t} + \sigma_\nu \sqrt{1-\phi_c^2} dB_{\nu,t} \\ d\xi_t &= \sigma_\xi \phi_x dB_{x,t} + \sigma_\xi \sqrt{1-\phi_x^2} dB_{\xi,t}. \end{aligned}$$

Following Epstein and Schneider (2007) and Illeditsch (2011), we assume that the agent is ambiguous about the precision of signals and thus considers a set of precisions when processing these signals,[IA.2]

$$\sigma_\theta^2 \in [\overline{\sigma_\theta}^2, \underline{\sigma_\theta}^2], \quad \theta = y, \nu, \xi. \tag{IA.A4}$$

Compared to their modeling framework, our model introduces an additional dimension of ambiguity; that is, the agent also feels uncertain about the relevance of (private) signals and perceives a range of correlation coefficients:

$$\phi_\theta \in [\overline{\phi_\theta}, \underline{\phi_\theta}] \subseteq [0,1], \quad \theta = c, x. \tag{IA.A5}$$

Let Υ denote the five-dimensional parameter vector, $\Upsilon = (\sigma_y^2, \sigma_\nu^2, \sigma_\xi^2, \phi_c, \phi_x)$, and let $\mathfrak{d}$ denote the convex set in $\mathscr{R}^5$ as defined by (IA.A4) and (IA.A5).

[IA.2] An alternative approach to modeling learning under ambiguity is described by Leippold et al. (2008). Their model focuses on ambiguous signals that are potentially biased in a nonparametric way.

In this model, we combine an ambiguous initial state with ambiguous signals. For example, the initial value of state variable x_0 is perceived to lie in the set $[-A_0, A_0]$. Empirically, the initial lack of confidence could be a reflection of dispersion in the time-0 forecasts of experts, whose opinions agents observe as signals. Utility at time 0 is given by

$$V_0^* = \min_{x_0 \in [-A_0, A_0]} E_{P(x_0)} \left[\int_0^\infty f\left(C_s, V_s^{P(x_0)}\right) ds \right]. \tag{IA.A6}$$

Given the aggregator function as presented in Eq. (5), it can be verified that the value function is nondecreasing in the conditional mean of consumption growth. It follows that $\hat{x}_0^* = -A_0$ supports the optimal choice in Eq. (IA.A6).

To construct the belief set $\{\hat{x}_t\}$ for $t > 0$, we update each priori belief by Bayes' rule. Epstein and Schneider (2003) show that, with multiple-priors utility, this prior-by-prior Bayesian updating is a necessary condition for the dynamic consistency of conditional preferences. If there is no ambiguity about signal quality, the updating rule for each prior can be obtained by applying the extended Kalman-Bucy filter:

$$\begin{bmatrix} d\hat{c}_t \\ d\hat{x}_t \end{bmatrix} = \begin{bmatrix} \bar{\mu} - \frac{1}{2}\sigma_c^2 + \hat{x}_t \\ -\kappa_x \hat{x}_t \end{bmatrix} dt + \begin{bmatrix} \frac{S_{cc,t}}{\sigma_y^2} & (\psi_c^2 + \psi_x^2 \kappa_x^2) S_{cx,t} \\ \frac{S_{cx,t}}{\sigma_y^2} & (\psi_c^2 + \psi_x^2 \kappa_x^2) S_{xx,t} \end{bmatrix} \begin{bmatrix} c_t - \hat{c}_t \\ x_t - \hat{x}_t \end{bmatrix} dt +$$

$$\begin{bmatrix} \frac{S_{cc,t}}{\sigma_y} & -\psi_c S_{cx,t} & \kappa_x \psi_x S_{cx,t} \\ \frac{S_{cx,t}}{\sigma_y} & -\psi_x S_{xx,t} & \kappa_x \psi_x S_{xx,t} \end{bmatrix} \begin{bmatrix} dB_{y,t} \\ dB_{\nu,t} \\ dB_{\xi,t} \end{bmatrix},$$

where $\psi_c = \frac{\phi_c}{\sigma_c \sqrt{1-\phi_c^2}}$ and $\psi_x = \frac{\phi_x}{\sigma_x \sqrt{1-\phi_x^2}}$.[IA.3] $S_{cc,t}$, $S_{cx,t}$, and $S_{xx,t}$ are elements of the conditional covariance matrix $S_t = E_t[(c_t - \hat{c}_t, x_t - \hat{x}_t)(c_t - \hat{c}_t, x_t - \hat{x}_t)'] \in \mathscr{R}^{2\times 2}$, which satisfies the following

[IA.3]We make a strong assumption about the structure of state processes corresponding to each prior. That is, all priors contained in the belief set use the same type of model to describe how consumption growth evolves, which coincides with the process presented in Eq. (IA.A1). This assumption is empirically relevant: even if all experts participating in the survey use the same model for forecasting, investors merely observe the forecasting results and still feel uncertain about how to model the dynamics of consumption growth. Theoretically, if the entire set of priors points to a single state process, the representative agent is arguably experiencing full confidence in this model structure. While this specification conceptually deviates the assumption made in Section 2.2 that the agent is blind to the true DGP, we find that the asset pricing implications derived from our benchmark model are robust even if we loosen this assumption. Specifically, expected consumption growth is assumed to truly follow an Ornstein-Uhlenbeck process as specified in Eq. (IA.A1); the agent acquires complete knowledge about its dynamics but still has to make inferences about its values. The key to the robustness of our baseline results is using the consensus survey forecast to form estimates of model parameters associated with x_t. That is, if we follow Bansal and Shaliastovich (2013) by using the consensus survey forecast to measure the true expected growth rate $\bar{\mu} + x_t$, then the MLE of the joint process for $\{C_t, x_t; O_t; \widetilde{A}_t\}$ reveals that (annualized) $\hat{\kappa}_x$ equals 0.56. As long as the x_t process is *modestly* persistent, the calibration results would show that the model performance is qualitatively the same as that presented in Section 4.2.1. Otherwise, if κ_x is calibrated to a small value as in long-run risk models (between 0.02 and 0.025), the main risk channels proposed in this paper would be obscured. More importantly, the model would overshoot the equity premium and underpredict the risk-free rate. In this paper, we aim to show how our model can capture empirical regularities in the stock and corporate debt markets with the simplest specification.

Riccati equation:

$$\frac{dS_t}{dt} = \begin{bmatrix} 2S_{cx,t} & S_{xx,t} - \kappa_x S_{cx,t} \\ S_{xx,t} - \kappa_x S_{cx,t} & -2\kappa_x S_{xx,t} \end{bmatrix} + \begin{bmatrix} S_{cc,t} & S_{cx,t} \\ S_{cx,t} & S_{xx,t} \end{bmatrix} \begin{bmatrix} \frac{1}{\sigma_y^2} & 0 \\ 0 & \psi_c^2 + \psi_x^2\kappa_x^2 \end{bmatrix} \begin{bmatrix} S_{cc,t} & S_{cx,t} \\ S_{cx,t} & S_{xx,t} \end{bmatrix} + \begin{bmatrix} \sigma_c^2 & 0 \\ 0 & \sigma_x^2 \end{bmatrix}.$$

If Knightian uncertainty only concerns the initial state of the economy, i.e., the entire dynamic system is perfectly understood and unchanging, the agent would eventually learn the true value of x_t and the belief set $\{\hat{x}_t\}$ would become degenerate in the limit (Epstein and Schneider, 2007). For the ambiguity to persist in the long run, at least some aspect of the signal-generating process needs to be ambiguous. Let us first consider the scenario of ambiguous signal precision as considered by Epstein and Schneider (2007) and Illeditsch (2011). It corresponds to the special case that ν_t and ξ_t are known to be irrelevant to consumption growth, i.e., $\phi_c = \phi_x = 0$. For a given prior $\hat{x}_t$ at time t, an application of Bayes' rule does not lead to a single posterior but rather a family of posteriors. To gain more intuition, we express the belief updating in discrete time:

$$\hat{x}_{t+1} = (1-\kappa_x)\hat{x}_t + \frac{(1-\kappa_x)^2(Q_t - \sigma_x^2)}{Q_t - \sigma_x^2 + (1-\kappa_x)^2(\sigma_c^2 + \sigma_y^2)} \left(\Delta y_{t+1} - \frac{\hat{x}_t}{1-\kappa_x} \right), \quad \sigma_y^2 \in [\underline{\sigma_y}^2, \overline{\sigma_y}^2].$$

It follows that the worst-case scenario is an imprecise signal ($\sigma_y^2 = \overline{\sigma_y}^2$) if good news arrives ($\Delta y_{t+1} > \frac{\hat{x}_t}{1-\kappa_x}$) and a reliable signal ($\sigma_y^2 = \underline{\sigma_y}^2$) if bad news comes in ($\Delta y_{t+1} < \frac{\hat{x}_t}{1-\kappa_x}$). This asymmetric response to ambiguous information implies that the worst-case posterior is more persistent than the single posterior in the case of known signal precision ($\sigma_y^2 \equiv \sigma_{y,0}^2$). Indeed, in the latter situation the mean-reverting parameter of $\hat{x}_t$ is greater than κ_x, the mean-reverting parameter of the true expected growth rate.

Besides signal precision, information quality in our model can be measured via the relevance of intangible signals as well. Consider another special case in which the agent considers the public signal y_t to be extremely noisy ($\sigma_y^2 \to \infty$). In this situation, it can be verified that one private signal ν_t becomes unusable regardless of its accuracy. Therefore, the agent can only rely on another private signal ξ_t for the inference of the expected growth rate:

$$\hat{x}_{t+1} \approx (1-\kappa_x)\hat{x}_t + \frac{\sigma_x \phi_x}{\sigma_\xi} \Delta\xi_{t+1}, \quad (\sigma_\xi^2, \phi_x) \in [\underline{\sigma_\xi}^2, \overline{\sigma_\xi}^2] \times [\underline{\phi_x}, \overline{\phi_x}].$$

As the agent is concerned about the worst-case scenario for the posterior conditional mean, she acts as if it is more likely that good news ($\Delta\xi_{t+1} > 0$) is generated by an imprecise and irrelevant signal ($\sigma_\xi^2 = \overline{\sigma_\xi}^2, \phi_x = \underline{\phi_x}$), whereas bad news ($\Delta\xi_{t+1} < 0$) is generated by a reliable signal ($\sigma_\xi^2 = \underline{\sigma_\xi}^2, \phi_x = \overline{\phi_x}$). In our simulation analysis, we find that incorporating ambiguous relevance of signals further increases the persistence of the ambiguity process.

IA.A.2. Parameterizations and Simulation Results

For model tractability, in this subsection we do not consider the interaction between ambiguous priors and ambiguous signals in the following discussion. Rather, each model contained in the set of beliefs corresponds to a fixed combination of the prior about the conditional growth rate x_t and parameters Υ in the signal-generating process. As such, we use $\mathcal{M}_t = \{x_{i,t}, \Upsilon_i : 1 \leq i \leq N\}$ to denote the set of all models considered by the agent at time t. This specification maps nicely onto the way ambiguity is empirically measured in Section 3.1: each of the N professional forecasters holds a distinct belief about the information quality Υ_i; they start from different priors about the initial state of the economy $\hat{x}_{i,0}$ and update their estimates of conditional consumption growth sequentially with the received signals and the forecaster-specific parameters Υ_i. The representative investor, on the other hand, surveys these experts for their forecasts of consumption growth in the next period. Following our definition in Section 2.2, the ambiguity level perceived by investors at time t can be captured by

$$A_{x,t} = (\max_{1\leq i\leq N} \hat{x}_{i,t} - \min_{1\leq i\leq N} \hat{x}_{i,t})/2. \tag{IA.A7}$$

With this specification, the ambiguity level tends to move slowly with signals. For instance, suppose that $i_t^* = \arg\min_i \hat{x}_{i,t}$ corresponds to the worst-case belief chosen from the set of posteriors at time t, and that a highly positive signal of ξ_{t+1} is received in the next period. Then it is likely that another prior j with moderately pessimistic expectation at time t but the most extreme values of $\sigma_{j,\xi}^2$ and $\phi_{j,x}$ would replace i_t^* as the worst-case belief at time $t+1$. Hence, the worst-case conditional mean would not move so fast as is the case of a single prior about x and Υ ($\mathcal{M}_t \equiv M_t$).

While the illustration as provided above suggests that the ambiguity level tends to follow a persistent stochastic process, we do not have a closed-form solution to the A_t defined in Eq. (IA.A7). In the following, we use Monte Carlo simulations to deduce how the model-implied ambiguity level varies over time. Especially, we aim to address two questions: (1) Is the exogenous ambiguity process specified in Eq. (10) a reasonable approximation of the endogenous A_t as derived from the learning model? (2) Does the endogenous A_t resemble the empirical measure $\widetilde{A}_t$ as introduced in Section 3.1, in terms of their magnitude and time-series properties? To directly answer the second question, we calibrate the true DGP to real output growth, which is the variable survey participants are asked to forecast.[IA.4] Specifically, parameter values in (IA.A1) are based on the maximum likelihood estimation with consensus forecasts (of real GDP growth from Blue Chip) as a proxy for x_t. Since the artificial sample is simulated at a quarterly frequency, we calibrate the public signal y_t to the second release of real GDP growth, which usually comes out during the third month of the next quarter. The historical second release data are taken from the "Real-Time Dataset for Macroeconomists" maintained by the Federal Reserve Bank of Philadelphia. Parameters involved in the DGP are summarized in Panel A of Table IA.A1.

We first study how close the conditional distribution of $A_{x,t}$ as implied by the learning model is to that of our empirical ambiguity measure, based on 5,000 finite-sample simulations. That is, the length of each simulation is 104 quarters, which is the length of the full sample (1985-2010) studied in this paper.[IA.5] An initial draw of the state variable x_t is taken from its stationary distribution. Subsequent draws of x_t and other variables use their conditional multivariate normal distribution.

Panel B summarizes the results derived from the finite-sample simulations. Columns (2) and (3) show that the average sample mean and sample variance across the 5,000 trials are about 0.90% and 0.62 bps, largely in line with their empirical counterparts — the first two moments of $\widetilde{A}_t/2$ — 1.03% and 0.73 bps. The next question is whether $A_{x,t}$ follows a stationary process varying around this mean in simulations. First, if the ambiguity will be resolved asymptotically, we could detect a deterministic time trend. To this end, we perform t-tests on the following stylized model:

$$A_{x,t+1} = \alpha + \beta t + \rho A_{x,t} + \varepsilon_{t+1}. \tag{IA.A8}$$

As reported in Column (4), the average value of estimated β is fairly small and not significant at any level. This result suggests that we exclude the time trend in unit-root tests.[IA.6] Column (5) shows

[IA.4] Based on the estimates of model parameters as listed in Table 2, the Knightian uncertainty in output growth can approximately be expressed as $0.94 \times A_t$. Hence, theoretically it should closely mimic the ambiguity in consumption growth in terms of both the absolute level and the temporal variation.

[IA.5] There is substantial evidence for a structural break in the early 1980s, following a significant policy experiment (Taylor, 1999; Clarida et al., 2000). Yang (2012) estimates a two-regime model on quarterly consumption growth in the period 1947-2012, and the estimated regime probabilities suggest that the economy may have switched into the "bad" regime (with low growth and high volatility) during the 1980 recession and then switched back to the "good" regime after 1982. Therefore, it is reasonable to assume that the agent considers observables from the pre-1985 period to be generated from a difference dynamic system and thus does not include them in her learning sample.

[IA.6] Unreported results show that the intercept term is highly significant in simulations. We also calculate the Bayesian information criterion (BIC) for different versions of Eq. (IA.A8) with or without the restrictions $\alpha = 0$ and $\beta = 0$

that the null hypothesis $\rho = 1$ is overwhelmingly rejected in the augmented Dickey-Fuller (ADF) test, and the Phillips-Perron test yields very similar p-values (unreported). Hence, we conclude that there is neither a deterministic trend nor a stochastic trend in the $A_{x,t}$ process.

To further explore the time-series properties of $A_{x,t}$, we calculate the first-order autocorrelation for each simulated sample. Column (6) shows that the average ACF(1) across simulations is 0.832. Accordingly, unreported results show that the average ρ in Eq. (IA.A8) (estimated with the constraint $\beta = 0$) is 0.911, which corresponds to a half-life of nearly 1.86 years. All the evidence indicates that $A_{x,t}$ follows a fairly persistent process, despite a much lower degree of persistence for the variation in individual priors' expected growth rate $\hat{x}_{i,t}$.[IA.7] This result stems from our definition of ambiguity — the gap between the most optimistic and the most pessimistic expectations — and from our model feature that these most extreme expectations may correspond to different priors over time.

The next few columns focus on the second moment of the ambiguity level. We firstly perform the heteroskedasticity test of Breusch and Pagan (1979)[IA.8] on the AR(1) model in Eq. (IA.A8), and Column (7) shows that the p-values average out at 0.005 over the 5,000 simulated samples. It follows that the variance of ambiguity shocks is dependent on the current level of ambiguity. Then, we test the direction of dependence by examining the slope coefficient in the Breusch-Pagan regression:

$$\hat{\varepsilon}_{t+1}^2 = \gamma_0 + \gamma_1 A_{x,t} + u_{t+1}.$$

As shown in Column (8), γ_1 is significantly positive, which is consistent with the direction of asymmetric volatility in our ambiguity measure $\widetilde{A}_t$. Hence, it is not surprising to find that the model-implied ambiguity level exhibits a positive skewness of 1.013 (Column (9)). Finally, we conduct the variance ratio test on $\hat{\varepsilon}_{t+1}$ as a robustness check on the time-varying ambiguity variance. While only the results with the horizons of 2 and 10 quarters are reported (in Columns (10) and (11)), variance ratios with other horizons show the same pattern: they are all above one and increase with the horizon.

Overall, the ambiguity process derived from our learning model justifies the CIR-type of speci-

(Hacker, 2010). Consistent with the conclusion hypothesis testing, the model with an intercept term but without a time trend has the minimum BIC.

[IA.7] It can be verified analytically that $\hat{x}_{i,t}$ is less persistent than the true expected growth rate x_t, of which the quarterly persistence is set at 0.731.

[IA.8] We implement the modified version of Koenker (1981) that is robust to non-normality.

fication in Eq. (10), which implies that the variation in aggregate ambiguity is driven by a mean-reverting component and by shocks with the conditional volatility varying positively with the ambiguity level. A major concern about the results in Panel B is whether the finite-sample properties of $\widetilde{A}_t$ remain stable in the long run. To answer this question, we look into its population distribution in Panel C. Specifically, artificial data of 5,000 quarters are generated for each trial, and we consider 5,000 trials as before.

Columns (2) and (3) shows that the population mean and variance are close to the corresponding values in Panel B. To test whether other aspects of the distribution function alter over time, we perform two-sample Kolmogorov-Smirnov tests on the 5,000 observations of $A_{x,t}$ for fixed ts. As shown in Column (4), the test fails to reject the null hypothesis that $A_{x,50}$ and $A_{x,100}$ have the same distribution function. And results in Columns (5)-(7) indicate that this conclusion holds for $t = 100$ vs. $t = 500$, $t = 500$ vs. $t = 1000$, and $t = 1000$ vs. $t = 5000$ as well. We also carry out a unified test to investigate whether the observations for these five different ts can be considered as arising from the same distribution, and the results are reported in Column (8). Again, we find that a five-sample Anderson-Darling test (Scholz and Stephens, 1987) cannot reject the null.

After establishing that the exogenous CIR process provides a good approximation of the learning-model-implied $A_{x,t}$, we finally try to validate our model parameterizations shown in Section 4.1. That is, if our model provides a reasonable description of the learning-under-ambiguity problem in the real world, then fitting the CIR model with the simulated $A_{x,t}$ should yield parameter values that are similar to those estimated from our survey data. Columns (9)-(11) report the average estimates of $\{\kappa_a, \bar{A}, \sigma_a\}$ across simulations; they are not significantly different from their counterparts in Table 2. Especially, the mean of the mean-reverting parameter κ_a, 0.295, is very close to the 0.303 based on professional forecasts, implying a half-life of about two years. Hence, the high persistence is a robust feature of the ambiguity process both in theory and in the data.

For robustness, we also perform simulation analysis for a slightly different model specification, which bears a closer resemblance to our empirical setting. On the one hand, experts may believe that they have learned the precision of the public signal y_t, after many observations of the BEA's non-final releases before 1985, and thus become certain about the true value of σ_y^2. On the other hand, it is plausible that the private signals received by experts differ from individual to individual. Therefore, they are assumed to be prior-specific and denoted by $\{\nu_{i,t}, \xi_{i,t}\}$ in our second simulation exercise.

The new simulation results are presented in Table IA.A2. Panel A displays the finite-sample properties of $A_{x,t}$ based on this variant specification of our learning model. A comparison with the corresponding results in Panel B of Table IA.A1 does not reveal any substantial difference. The only noteworthy change is that both the overall level and the volatility of $A_{x,t}$ slightly increase, edging closer toward their empirical counterparts. Similarly, the population properties shown in Panel C of Table IA.A1 persist in the alternative model specification. Again, the stationary moments of $A_{x,t}$ and the estimates of its dynamics presented in Panel B of Table IA.A2 seem to show a closer similarity to those derived from the survey data. This finding suggests that forecaster-specific private signals can be considered as a justifiable assumption in the real world.

Since our model takes into account the empirical consideration that the growth rate realized at time t is not immediately observable, the agent at time t feels uncertain about it as well. For this reason, in the first publication of each quarter (e.g., the issue published on January 1, 2000), the BCFF includes in the "Special Questions" section the participants' responses to a question about their expectations of real GDP growth realized in the last quarter (e.g., the growth rate corresponding to the fourth quarter of 1999).[IA.9] The dispersion in these responses provides another setting in which our model can be tested against the data. That is, in our model, forecaster i's time-t projection of output growth realized at time t can be expressed as

$$\Delta\hat{c}_{i,t} = \bar{\mu} - \frac{1}{2}\sigma_c^2 + \frac{\hat{x}_{i,t}}{1-\kappa_x} - \frac{\sigma_x\phi_x}{\sigma_\xi(1-\kappa_x)}\Delta\xi_{i,t} + \frac{\sigma_c\phi_c}{\sigma_\nu}\Delta\nu_{i,t}. \quad \text{(IA.A9)}$$

Given that our point estimate of $\hat{\kappa}_x = 0.28$[IA.10] is reasonable, the dispersion in contemporaneous forecasts is on average greater than but close to that in one-period-ahead forecasts. This is consistent with what the survey data reveals: if our ambiguity measure for expected growth (the interdecile range) is applied to contemporaneous forecasts, the resultant historical average is 2.18%,

[IA.9] This question is asked in the January, April, July, and October issues only. For a few of these issues where it is missing, we collect the corresponding data from the SPF, which always contains the forecasters' projections for the current quarter. The main problem with these so-called contemporaneous forecasts is about the timing of surveys; that is, SPF questionnaires are sent during the first week of the middle month of each quarter, and the deadlines are usually set in the second week. As a consequence, panelists in the SPF survey generally face greater uncertainty when making forecasts of the growth rate of the current quarter, compared to BCFF forecasters. To alleviate this problem, we perform a linear interpolation between a participant's "zero-quarter-ahead forecast" (the "RGDP2" item) in quarter i's survey and her "one-quarter-behind forecast" (the "RGDP1" item) in quarter $i+1$'s survey, given that she participates in both surveys, to construct a "true" contemporaneous forecast. To be more specific, the forecasting horizon is 1.5 to 2 months for the former and -1.5 to -1 months for the latter, so that we can obtain a "forecast" with a horizon of almost 0 months by interpolation.

[IA.10] This estimate corresponds to the annualized κ_x of 0.72 in Panel A.

which is slightly greater than the 2.07% for the measure for one-quarter-ahead forecasts. As suggested by Eq. (IA.A9), this difference mainly results from the missing role of y_t in the inference of Δc_t, as the public signal is assumed to have a one-period lag and thus is only informative about future growth rates.

To examine the quantitative implications for the uncertainty in contemporaneous growth rates, we simulate $\{\Delta \hat{c}_{i,t}, 1 \leq i \leq N\}$ from the same DGP as in Panels A and B and compute the moments of their cross-sectional range

$$A_{c,t} = (\max_{1 \leq i \leq N} \Delta \hat{c}_{i,t} - \min_{1 \leq i \leq N} \Delta \hat{c}_{i,t})/2.$$

Column (1) in Panel C shows that the finite-sample mean of $A_{c,t}$ is 1.06%, which is extremely close to its empirical counterpart of 1.09% as reported in Column (9). As expected, it is slightly greater than the mean of $A_{x,t}$ (0.95%). This result holds for the population mean as well (1.06% against 0.98%). On the other hand, the variance of $A_{c,t}$ in both small-sample and long-sample simulations (Columns (2) and (6)) is somewhat lower than that of $A_{x,t}$. The survey data show the same pattern: the interdecile range of one-quarter-ahead forecasts is more volatile than that of contemporaneous forecasts, as their variances are 2.77 bps and 2.64 bps (four times the 0.66 pbs shown in Column (10)), respectively. Another property of interest is the skewness of $A_{c,t}$. Column (11) indicates that the dispersion in forecasters' contemporaneous projections is highly right-skewed. The learning model captures this positive skew as well. Finally, both the survey data and model simulation show that the uncertainty level of realized output growth is less persistent than that of expected output growth.

Appendix IA.B. Ambiguity about Real GDP Growth

IA.B.1. Details of the Data Set

The BCFF dataset provides an extensive panel of data on expectations by professional economists from leading financial institutions and service companies. Each economist is asked to forecast a large set of macroeconomic and financial variables. Our ambiguity measure draws on the individual-level forecasts of real gross national product (GNP). The name of this variable was changed to real gross domestic product (GDP) after January 1992.

Besides the Blue Chip survey, the Survey of Professional Forecasters (SPF) is another widely used source of survey data.[IA.11] For the objective of this study, the BCFF data appear to bear greater relevance, for three reasons.

First, as shown in Table IA.B3, the BCFF has kept the number of respondents fairly stable over the last 30 years, which is always within the interval between 44 and 50. By contrast, the number of forecasters participating in the SPF varies enormously, from 9 to 53. This feature is undesirable, at least in the multiple-priors framework, because changes in the subjective set of priors should not be a mere reflection of rapid turnover of survey participants. Instead, they are supposed to mirror updates of beliefs.

Second, Table IA.B3 indicates that the composition of the BCFF's forecaster panel remains almost constant over time. On average, about 94% of analysts appearing in the current period's publication also participate in the survey of the preceding period. By contrast, it is not unusual that a batch of respondents suddenly drop out of the SPF for a large number of periods and then suddenly reenter.[IA.12] If the forecast is more closely associated with the individual participant, the highly fluctuating composition makes the panel unsuitable for our extraction of the set of priors.

Third, the BCFF survey is carried out within a short window toward the end of each month (between the 25th and 27th) and mailed to subscribers within the first five days of the subsequent month. This allows the expectation horizon of panelists to closely match our sample frequency and ensures that our analysis of return predictability in Section 3.2 does not involve overlapping observations (between returns and ambiguity measures). By comparison, participants in the SPF receive questionnaires by the end of the first month of each quarter and are asked to complete them within two weeks (sometimes three weeks in surveys before 2005:Q1).[IA.13] As a result, there are no effective one-quarter-ahead forecasts in the SPF data.

IA.B.2. Distribution of the Proposed Ambiguity Measure

In Section 3.1, we construct a proxy $\widetilde{A}_t$ for aggregate ambiguity. As suggested by its historical variation as charted in Figure 1, this ambiguity measure seems to exhibit heteroskedasticity and a positive skew in its distribution. We further explore these empirical properties in this subsection by

[IA.11] The Livingston Survey also provides economists' forecasts of real GDP growth rate, but at a semiannual frequency.

[IA.12] The problems associated with the changing composition of the SPF panel are discussed exhaustively in Engelberg et al. (2011).

[IA.13] This information only applies to surveys conducted by the Federal Reserve Bank of Philadelphia, that is, surveys after 1990:Q2. The timing of surveys in the ASA/NBER era is unclear.

first performing a variance ratio test on the realized volatility of innovations. Panel A of Table IA.B4 displays the test results. If realized volatility were i.i.d., the term structure of variance ratios would be flat, with magnitudes close to unity. But in Table IA.B4 the empirical variance ratios are all above one, and they increase uniformly with the horizon. The bootstrap analysis indicates that the null of homoskedastic shocks is rejected at the 5% level for all horizons.

Even more prominent is the skewness of the ambiguity level. As shown in Panel B, the number measured for our sample period is 0.82. This positive skewness is significant at the 5% confidence level constructed by 10,000 bootstrap samples. This evidence, together with the heterogeneity of realized volatility that is positively correlated with A_t, prompts a square-root specification for the diffusion term in the A_t process:

$$dA_t - E(dA_t|\mathscr{F}_t) = \sqrt{a_0 + a_1 A_t} dB_{A,t}, \quad a_1 > 0.$$

Since we also need to guard against the probability that A_t falls below zero, we restrict a_0 to zero and thus adopt the CIR specification as shown in Eq. (10). Figure IA.B1 illustrates that the density function implied by the simulated CIR process fits the empirical one reasonably well, while a simple Ornstein–Uhlenbeck process, with $dA_t - E(dA_t|\mathscr{F}_t) = \sigma_a dB_{A,t}$ for a fixed σ_a, is unable to capture the right skewness.

It is worth noting that we have also calculated the skewness and variance ratios for each simulated sample in Section IA.A.2, and their average values as shown in Tables IA.A1 and IA.A2 appear close to their empirical counterparts. Indeed, unreported results show that the sample statistics presented in Table IA.B4 are all within the 95% confidence intervals as implied by the learning model.[IA.14] Therefore, the dynamic properties of our ambiguity measure not only justify the exogenous CIR process adopted to derive theoretical implications for asset pricing, but also share many similarities with the ambiguity process implied from a structural model of learning.

IA.B.3. The Predictive Ability of $\widetilde{A}_t$ Conditional on Additional Predictor Variables

Results in Section 3.2 show that (a) $\widetilde{A}_t$ has short-horizon forecasting ability in univariate regressions, and (b) this finding is robust to the inclusion of predictor variables that have been used in earlier works to forecast asset returns at horizons of one month and/or one quarter. Although it is well

[IA.14] These confidence intervals are constructed based on the 5,000 simulated data sets.

documented in the literature that forecasting short-run returns is more challenging than forecasting long-run ones,[IA.15] at least in in-sample analysis, we stick to the one-quarter horizon, as our model implication (Eq. (24)) is about the impact of ambiguity on the expected excess returns in the next period. With that being said, one cannot preclude the possibility that the predictive power of $\widetilde{A}_t$ partly stems from its correlation with those long-horizon predictor variables. In this section, we examine the significance of $\widetilde{A}_t$ once these variables are controlled for in predictive regressions.

To be more specific, we consider three predictor variables shown to be significant in forecasting equity returns at horizons of one year or longer. The first variable is the net payout yield, where the numerator is defined as dividends plus repurchases minus issuances; Boudoukh et al. (2007) document that it is a robust predictor for one-year-ahead stock returns and substantially outperforms the dividend yield. The second variable is the detrended three-month T-bill rate, which has been shown by Fama and Schwert (1977) to be able to predict the equity return. Finally, we include Lettau and Ludvigson (2001)'s consumption-to-wealth ratio *cay*.

Table IA.B5 examines whether the predictability with $\widetilde{A}_t$ is robust to the inclusion of this set of alternative predictor variables. In-sample regressions indicate that $\widetilde{A}_t$ retains its forecasting power with roughly the same coefficient size and same level of statistical significance. Judging from the R^2 values, these long-horizon predictors contain little information about the short-run stock returns exceeding that of $\widetilde{A}_t$. The out-of-sample results point to the same conclusion.

Then we examine the robustness of the predictability of credit spreads. Due to the lack of well-recognized spread predictors, we consider three (aggregate) variables that are significant in explaining contemporaneous credit spreads (Collin-Dufresne, Goldstein, and Martin, 2001): changes in the 10-year T-note rate (a proxy for changes in the Treasury yield level), changes in the VIX index (a proxy for changes in aggregate volatility), and S&P 500 returns (a proxy for the expected recovery rate). To save space, we do not report results by alphanumeric rating, but focus on the yield spreads of two broad categories in Panel B: investment-grade and speculative-grade bonds. We find that the forecasting ability of $\widetilde{A}_t$ is not affected by the inclusion of control variables. The significance of these control variables is limited to high-yield bonds.

Finally, we revisit the predictive power of $\widetilde{A}_t$ for Treasury bond returns by controlling for two

[IA.15] Regarding equity returns, Cooper and Priestley (2009) and Jones and Tuzel (2013) find that with the same set of predictors, the R^2 value strictly rises with the horizon. Duffee (2007) and Mueller et al. (2011) present similar evidence for Treasury bond returns.

variables that can achieve a high degree of predictability for annual excess bond returns. Specifically, Cochrane and Piazzesi (2005) find that a tent-shaped linear combination of forward rates produces R^2 values ranging from 30% to 35%, and it greatly outperforms the single forward spread used in Fama and Bliss (1987). Ludvigson and Ng (2009) summarize information in 132 macro-finance series into a single factor, which can explain 21%-26% of variation in next year's excess bond returns. In conjunction with these two factors, $\widetilde{A}_t$ remains statistically significant at the 5% level in all but the last column in Panel C. That is, the p-value of $\widetilde{A}_t$ is 5.26% in the *in-sample* regression of long-term bond returns on the Ludvigson-Ng factor along with it. It suggests the possibility that the information in our ambiguity measure is partially absorbed by the large panel of macroeconomic data used in Ludvigson and Ng (2009). However, the Clark-West test indicates that $\widetilde{A}_t$ makes a significant contribution to the *out-of-sample* forecast when the Ludvigson-Ng factor is controlled for, with a p-value of 2.89%.

Appendix IA.C. Market-Based Estimation of the Default Boundary and the Cost of Default

IA.C.1. Data Description

As is typical in studies of the default cost and recovery rates, our estimation relies on several data sources. We start with the master list of defaults provided by the Altman-Kuehne NYU Salomon Center; it contains all public, nonconvertible corporate debt issues that have either filed for bankruptcy or defaulted on a scheduled interest or principal payment since 1986. Driven primarily by data availability, the focus on the database underlying the Altman-Kuehne Defaulted Bond Index implies that many (out-of-court) distress renegotiations are not included in our analysis.

As shown in Panel A of Table IA.C6, the overall sample consists of 1,479 defaulted bonds issued by 696 firms.[IA.16] Defaults appear to cluster during the early 1990s recession, the dot-com crash of the early 2000s, and during the 2008 financial crisis (Giesecke et al., 2011). We merge defaulting firms with CRSP data to get month-end share prices, and this reduces our sample to 219 public defaulting firms. Then we search for bond prices from the Warga and Merrill Lynch databases as well as from Bloomberg Professional. Bond prices are available in the month prior to default for

[IA.16]If multiple default events occur within a firm or its wholly owned subsidiaries, we focus on defaults by the parent company only and look at the first default event during our sample period.

196 out of the 219 firms. Finally, post-default bond prices are extracted from the Altman-Kuehne database.

We simplify the estimation procedures of Davydenko et al. (2012, DSZ hereinafter) by estimating the market prices of bank loans as a linear-quadratic function of bond prices,

$$P_{bank} = 0.4018 + 1.045 \times P_{bond} - 0.461 \times P_{bond}^2, \tag{IA.C10}$$

where P_{bank} and P_{bond} are the average market prices of bank loans and public debts with a \$1 face value. The coefficients in Eq. (IA.C10) are taken from Appendix 2 of DSZ, who find that this specification can explain more than 75% of variations in bank debt prices available in their dataset. The total loan value, D_{bank}, is calculated as P_{bank} times the face value of bank loans; the latter is derived from the debt structure data of CapitalIQ.

IA.C.2. Properties of Market-Based Estimates

Our estimation of the default cost is based on its relationship with the firm-level price reaction to the default announcement and the risk-neutral default probability

$$E_{\tau_-} + D_{\tau_-} - (E_\tau + D_\tau) = \phi \times E^Q_{\tau_-}\left[\frac{U_T e^{-r(T-\tau)}}{U_{\tau_-}} e^{-\int_{\tau_-}^{T} \lambda_u^Q du}\right], \tag{IA.C11}$$

where ϕ, the deadweight value loss, and U_{τ_-}, the unlevered asset value at the time of default, are the focus of our analysis. Empirically, E_{τ_-} and $D_{\tau_-} (= D_{bank,\tau_-} + D_{bond,\tau_-})$ correspond to the firm's equity and debt values measured at the end of the calendar month preceding default, and E_τ and D_τ are their counterparts at the end of the calendar month of default. The risk-neutral default intensity, λ_t^Q, is defined as the physical-measure intensity λ_t^P times the jump-to-default risk premium. The former is estimated based on a hazard model (Shumway, 2001), while the latter is assumed to be constant over time. Its value is set at 2.31, which is the empirical estimate reported by Driessen (2005).[IA.17] DSZ find that their estimates are fairly insensitive to the assumption of this risk premium factor, whether it is specified as time-varying or fixed.

Panel B of Table IA.C6 summarizes our estimation results. The first few rows describe the market reaction in the month of default. The average firm-level monthly return is 15.9%, which is not significantly different from the -12.2% reported in DSZ. On the other hand, asset returns

[IA.17]Similar results are presented in Berndt et al. (2005), who estimate the jump-risk premium from the data on credit default swaps.

show marked variation across defaulting firms, ranging from -34.9% to +2.0% between the first and last deciles. Looking into the returns on specific asset classes, we find that for an average firm, equityholders tend to suffer heavier losses than debtholders (-33.6% vs. -13.4%). This reflects that debtholders generally have a senior claim on the firm's assets in bankruptcy compared to equityholders. Consistent with this interpretation, unreported results show that the mean return on bank loans is higher than that on public debts in the month of default. Finally, the overall debt recovery rates average 45.6%, lying between the issuer-weighted average recovery rate of senior unsecured loans (47.8%) and that of senior unsecured bonds (36.7%) as reported by Moody's.[IA.18]

The cost of default shown in Panel B is scaled by the market value of unlevered equity prior to default (c/U_{τ_-}). The mean of our estimates, 24.7%, is consistent with the 21.7% reported by DSZ, though they appear more homogeneous than the DSZ estimates.[IA.19] This result is actually to be expected, as we do not sample secondary market loan prices and instead approximate them using Eq. (IA.C10).

The last row of Panel B presents statistics for the solvency boundary, expressed as a proportion of the face value of total debt. We find that on average, the location of the default boundary is measured at 67.6% of the book value of debt, which is very close to the 66.0% reported by Reisz and Perlich (2007) and Davydenko (2012). It is also highly dispersed across firms, varying from 28.0% at the first decile to 104.8% at the tenth decile. Results in Section 4.2.3 show that the ambiguity level has strong explanatory power for its variation.

Appendix IA.D. Model Solutions Based on Asymptotic Expansions

The PDE characterization of the optimal stopping problem as presented in Appendix C resembles the problem confronted in pricing American options. For our asset pricing model where the risk-free rate and asset volatility are time-varying and dependent on the ambiguity level, two numerical approaches might be relevant: the short-maturity asymptotic approximation (Lamberton and Villeneuve, 2003; Chevalier, 2005; Levendorskii, 2008) and the simulation-based method proposed in Longstaff and Schwartz (2001). But neither approach is applicable here because there is no such

[IA.18] These statistics are based on the 1983-2010 sample period. See Ou et al. (2011) for details.

[IA.19] For example, the standard deviation of our estimates is 25.9% and that of DSZ estimates is 33.0%, and the interdecile range is 60.4% against 88.1%.

expiration date in the current debt structure. Therefore, we employ perturbation techniques to construct approximate asymptotic series expansions of contingent claim valuations:

$$E = E_0 + \sqrt{\varepsilon}E_1 + \varepsilon E_2 + O(\varepsilon^{\frac{3}{2}}), \tag{IA.D12}$$

$$D = D_0 + \sqrt{\varepsilon}D_1 + \varepsilon D_2 + O(\varepsilon^{\frac{3}{2}}). \tag{IA.D13}$$

These asymptotic approximations can be viewed as an extension of the slow variation asymptotics (Sircar and Papanicolaou, 1999; Lee, 2001) to optimal stopping-time problems (Fouque et al., 2001).

Given the scale parameter ε incorporated in Eq. (19), it seems convenient to write the operator as a sum of components that are separated by the different powers of ε:

$$\mathscr{L}_{\delta,A}^r = \mathscr{L}_0^r + \sqrt{\varepsilon}\mathscr{L}_1 + \varepsilon\mathscr{L}_2,$$

where

$$\begin{aligned}
\mathscr{L}_0^r &= \mu_\delta^Q \delta \frac{\partial}{\partial \delta} + \frac{1}{2}\bar{\sigma}^2\delta^2\frac{\partial^2}{\partial \delta^2} - r\cdot, \\
\mathscr{L}_1 &= \sigma_{oa}\sigma_a^2 A \frac{\partial^2}{\partial\delta\partial A} - [\gamma\sigma_{ca} + (1-\theta)\eta_1]\sigma_a^2 A \frac{\partial}{\partial A}, \\
\mathscr{L}_2 &= \kappa(\bar{A} - A)\frac{\partial}{\partial A} + \frac{1}{2}(\sigma_a^2 A)\frac{\partial^2}{\partial A^2}.
\end{aligned}$$

We note that $\mathscr{L}_0$ is simply the Black-Scholes operator with the risk-neutral drift

$$\mu_\delta^Q = \mu_c - \sigma_o\sigma_c\sigma_{oc}\gamma - \sigma_o\sigma_{oc}A/\sigma_c - [\gamma\sigma_{ca} + (1-\theta)\eta_1]\sigma_{oa}\sigma_a^2$$

and conditional variance

$$\bar{\sigma}^2 = \sigma_o^2 + \sigma_{oa}^2\sigma_a^2 A + \sigma_j^2.$$

$\mathscr{L}_2$ is the Dynkin operator of the A_t process under the physical measure. $\mathscr{L}_1$ contains the mixed derivative, due to the covariation between aggregate output and ambiguity, and the first-order derivative due to the uncertainty premium imposed on the A_t process.

As discussed in Section 2.4, we look for an asymptotic solution for E and D as an expansion in powers of $\sqrt{\varepsilon}$. The free boundary δ^* also has to be determined as part of the problem:

$$\mathscr{L}_{\delta,A}^r E + (1-\tau_e)(\delta - C) - mF + D(\delta, A) = 0, \tag{IA.D14}$$

$$\mathscr{L}_{\delta,A}^{r+m} D + (1-\tau_i)mC + m^2F = 0, \tag{IA.D15}$$

with boundary conditions given in Appendix C. With the log-linear approximation of the price-earnings ratio $L^o(A_t)$ and the perpetuity factor $\mathscr{P}\mathscr{V}^r(A)$, we obtain the following expressions for the leading-order terms:

$$
\begin{aligned}
D_0 &= ((1-\tau_i)mC + m^2F)e^{\zeta_0^m + \zeta_1^m A}\left(1 - \left(\frac{\delta}{\delta_0^*(A)}\right)^{\alpha_{r+m}}\right) + \\
&\quad m(1-\phi)\delta_0^*(A)e^{\xi_0+\xi_1 A}\left(\frac{\delta}{\delta_0^*(A)}\right)^{\alpha_{r+m}}, \qquad \text{(IA.D16)}\\
E_0 &= (1-\tau_e)\delta e^{\xi_0+\xi_1 A} + (\tau_e - \tau_i)Ce^{\zeta_0+\zeta_1 A}\left(1 - \left(\frac{\delta}{\delta_0^*(A)}\right)^{\alpha_r}\right) - \\
&\quad \phi\delta_0^*(A)e^{\xi_0+\xi_1 A}\left(\frac{\delta}{\delta_0^*(A)}\right)^{\alpha_r} - \frac{D_0(\delta, A)}{m}. \qquad \text{(IA.D17)}
\end{aligned}
$$

Inserting the expansion series in (IA.D12)-(IA.D13) into the PDE (IA.D14) and (IA.D15), we obtain

$$
\begin{aligned}
(1-\tau_i)mC + m^2F + \quad & \mathscr{L}_0^{r+m}D_0 + \sqrt{\varepsilon}(\mathscr{L}_1 D_0 + \mathscr{L}_0^{r+m}D_1) \\
+ \quad & \varepsilon(\mathscr{L}_2 D_0 + \mathscr{L}_1 D_1 + \mathscr{L}_0^{r+m}D_2) + O(\varepsilon^{\frac{3}{2}}) \quad = 0, \qquad \text{(IA.D18)}\\
(1-\tau_e)(\delta - C) - mF + \quad & D(\delta, A) + \mathscr{L}_0^r E_0 + \sqrt{\varepsilon}(\mathscr{L}_1 E_0 + \mathscr{L}_0^r E_1) \\
+ \quad & \varepsilon(\mathscr{L}_2 E_0 + \mathscr{L}_1 E_1 + \mathscr{L}_0^r E_2) + O(\varepsilon^{\frac{3}{2}}) \quad = 0. \qquad \text{(IA.D19)}
\end{aligned}
$$

Equating to zero the first few terms independent of ε leads us to the solution for the lead-order prices as shown in Eq. (IA.D16) and (IA.D17), where

$$
\begin{aligned}
\alpha_r &= -\frac{\mu_\delta^Q - \bar{\sigma}^2/2 + \sqrt{(\mu_\delta^Q - (\sigma_o^2 + \sigma_{oa}^2\sigma_a^2 A + \sigma_j^2)/2)^2 + 2r(\sigma_o^2 + \sigma_{oa}^2\sigma_a^2 A + \sigma_j^2)}}{\sigma_o^2 + \sigma_{oa}^2\sigma_a^2 A + \sigma_j^2}, \\
\alpha_{r+m} &= -\frac{\mu_\delta^Q - (\sigma_o^2 + \sigma_j^2)/2 + \sqrt{(\mu_\delta^Q - (\sigma_o^2 + \sigma_j^2)/2)^2 + 2(r+m)(\sigma_o^2 + \sigma_j^2)}}{\sigma_o^2 + \sigma_j^2}.
\end{aligned}
$$

Next, the order $\sqrt{\varepsilon}$ terms in the expansion (IA.D18) and (IA.D19) give the following equations for the first-order corrections:

$$
\begin{aligned}
\mathscr{L}_1 D_0 + \mathscr{L}_0^{r+m}D_1 &= 0, \\
\mathscr{L}_1 E_0 + \mathscr{L}_0^r E_1 &= 0.
\end{aligned}
$$

Boundary conditions are constructed by substituting the expansion series into original conditions

(27) and (31)

$$
\begin{aligned}
D_1(\delta_0^*) + \delta_1^* \frac{\partial D_0}{\partial \delta}|_{\delta=\delta_0^*} &= (1-\phi) m \delta_1^* e^{\xi_0+\xi_1 A_t}, \\
\lim_{\delta\to\infty} D_1(\delta) &= 0, \\
E_1(\delta_0^*) &= 0,
\end{aligned}
$$

where we have used the fact that $\lim_{\delta\to\infty} D_0(\delta)$ does not depend on δ. Finally, we can use (28) to find the order $\sqrt{\varepsilon}$ term in the expansion (21) for the optimal default barrier:

$$
\frac{\partial E_1}{\partial \delta}|_{\delta=\delta_0^*} + \delta_1^* \frac{\partial^2 E_0}{\partial \delta^2}|_{\delta=\delta_0^*} = 0.
$$

Since δ_0^* has been determined in Eq. (22), the first-order correction terms E_1 and D_1 actually solve a one-dimensional fixed boundary problem.

Likewise, the second-order price corrections are the solution to the following problem:

$$
\begin{aligned}
\mathscr{L}_2 D_0 + \mathscr{L}_1 D_1 + \mathscr{L}_0^{r+m} D_2 &= 0, \\
\mathscr{L}_2 E_0 + \mathscr{L}_1 E_1 + \mathscr{L}_0^{r} E_2 &= 0, \\
D_2(\delta_0^*) + \delta_1^* \frac{\partial D_1}{\partial \delta}|_{\delta=\delta_0^*} + \delta_2^* \frac{\partial D_0}{\partial \delta}|_{\delta=\delta_0^*} + \delta_1^{*2} \frac{\partial^2 D_0}{\partial \delta^2}|_{\delta=\delta_0^*} &= (1-\phi) m \delta_2^* e^{\xi_0+\xi_1 A_t}, \\
\lim_{\delta\to\infty} D_2(\delta) &= 0, \\
E_2(\delta_0^*) &= 0, \\
\frac{\partial E_2}{\partial \delta}|_{\delta=\delta_0^*} + \delta_1^* \frac{\partial^2 E_1}{\partial \delta^2}|_{\delta=\delta_0^*} + \delta_2^* \frac{\partial^2 E_0}{\partial \delta^2}|_{\delta=\delta_0^*} + \delta_1^{*2} \frac{\partial^3 E_0}{\partial \delta^3}|_{\delta=\delta_0^*} &= 0.
\end{aligned}
$$

Appendix IA.E. Convexity Bias in Credit Spreads

In many studies on the credit spread controversy, the spread calculations are carried out by averaging firm variables across firms within a rating category and over a historical time period. David (2008) argues that this "representative firm" calculation could lead to a large downward bias as credit spreads are a convex function of the solvency ratio. BKS (2010) further illustrate the importance of convexity by comparing a representative firm with a cross section of firms in their structural model.[IA.20] It is worth noting that their model-implied cross-sectional average is obtained by simulating a dynamic economy (which is populated by hundreds of firms) 1,000 times. Therefore, while

[IA.20] Note that BKS (2010) embed a structural credit model inside a representative-agent consumption-based model. While their results show that considering cross-sectional firm dynamics and heterogeneity increases the model spreads, the model would still be unable to account for the observed Baa-Aaa spreads without the proposed intertemporal macroeconomic risk.

this approach can fully address convexity bias, it can be computationally burdensome when used to test structural models. In this paper, we propose a "medians-to-medians" method that will be shown to be not susceptible to convexity bias. That is, we use median firm-level variables as inputs to the model and then compare the resulting model spread with the median actual spread.

In this appendix, we first explore the empirical difference between the median of cross-sectional spreads and spread indices, as the latter are (weighted) averages of cross-sectional yield spreads and are typically used as a reference for comparison with model-predicted spreads. To obtain a long-span bond sample covering the entire 1985-2010 period, we consolidate corporate bond data from Lehman Brothers and Merrill Lynch databases: for the 1985-1997 period, we extract all quotes (matrix prices are discarded) on noncallable US bonds from Lehman Brothers Fixed Income Database, whose availability is restricted after 1998; thenceforth, we switch to Merrill Lynch as the data source for bid quotes.[IA.21] To be included in the sample, a bond must belong to a Lehman (now Barclays) or Merrill Lynch (now BofA ML) index and have at least 24 *actual* monthly quotes during a 7- to 15-year maturity period. To construct credit spreads, we follow the market convention to use the nearest-maturity Treasury bond as the benchmark. Monthly yields of individual Treasuries are obtained from CRSP monthly bonds master file.

Table D1 shows that for each rating group, the expected spread for a typical bond is smaller than the average spread for the corresponding Barclays corporate bond index.[IA.22] Three factors might account for this spread discrepancy. First, the distribution of the credit spread is known to be right-skewed, and so the mean is likely to be higher than the median. Second, the bond pool underlying these indices includes callable bonds, which are known to have generally higher credit spreads than noncallable ones (Caouette et al., 1998). Third, our bond sample contains only bonds close to 10-year maturity, but the indices are based on a wide array of bonds with diverse maturities (their average time to maturity is 10.45 years). The second panel indicates that this discrepancy is not due to the switch to Merrill Lynch as the data source, as it has the same magnitude in the subsample before 1998. The last panel shows that over the second subsample period, the average

[IA.21] The reason for retrieving bid prices instead of mid prices is that the Warga Database contains only bid prices. Both data sources are widely used in the corporate bond literature (Collin-Dufresne et al., 2001; Elton et al., 2001; Eom et al., 2004; Chen et al., 2007; Schaefer and Strebulaev, 2008). We examine their comparability in terms of data quality and find a close match in their prices between December 1996 and December 1997, the period when these two datasets overlap.

[IA.22] Note that this is the index used by Huang and Huang (2012) as a reference for comparison with model-predicted 10-year spreads.

level of Barclays indices is uniformly higher than that of Merrill Lynch US Corporate/High Yield 7-10Y option-adjusted spreads; this finding lends support to the second interpretation above.

Next, we empirically assess the severity of the convexity bias in the conventional Merton model and then evaluate the medians-to-medians approach as a remedy for this bias. Panel A of Table D2 compares the "representative firm" credit spreads with the average model spreads, where the representative firm is defined as a firm with a leverage ratio, payout ratio, and asset volatility matching the historical averages in a particular rating category. An early version of Feldhütter and Schaefer (2018) performs a similar comparison for the 2002-2012 sample period.[IA.23] Consistent with their results, we find that convexity bias with the averages-to-averages approach is nontrivial and increases with credit quality. For 10-year Aaa bonds, the "representative firm" credit spread accounts for less than 40% of the average model spread; but for 10-year speculative-grade bonds, these two calculation methods result in the same magnitude of model spreads. In addition, the convexity bias seems especially marked for short-maturity bonds. The model spreads of a representative firm are almost negligible for one-year investment-grade bonds and less than 10% of the corresponding average spreads over a cross section of firms. However, entries in angle brackets show that no matter how the model spread is calculated, it always falls short of the observed yield spread.

Will the convexity bias persist if the representative firm is redefined by the median values of firm-specific variables in a rating category? Results in Panel B of Table D2 provide the answer. First, the empirical distributions of both the leverage ratio and asset volatility are right-tailed. On the one hand, this pattern suggests that the credit spreads of the newly defined representative firm tend to be lower than their counterparts in Panel A. On the other hand, the distribution of model spreads of individual bonds would become even more skewed compared to the case in which firm variables are symmetrically distributed. Put differently, the convex function for model spreads magnifies the positive skewness in model inputs. Consequently, when the means are replaced with medians, the reduction in the rows labeled "Model_CS" is larger than that in the "Model_Rep" rows. As a result, we do not find a significant difference in spread values from these two calculation

[IA.23] In order to draw a direct comparison with the median spreads, we do not weight bonds with outstanding volume, as Feldhütter and Schaefer (2018) do when taking the cross-sectional average. Another notable difference from their procedures is that we estimate asset volatility for individual firms using the KMV method (Vassalou and Xing, 2004; Crosbie and Bohn, 2003; Duffie et al., 2007).

methods in Panel B. Finally, the distribution of observed spreads also seems right-skewed, such that numbers in the "Actual" rows are uniformly smaller than their counterparts in Panel A. Nevertheless, they are still substantially greater than model spreads regardless of the calculation method. Overall, we demonstrate that the medians-to-medians comparison is generally immune to convexity bias and thus offers a convenient way to test models' predictions of credit spreads.

References

Bansal, R., Shaliastovich, I., 2013. A long-run risks explanation of predictability puzzles in bond and currency markets. Review of Financial Studies 26, 1–33.

Berndt, A., Douglas, R., Duffie, J., Ferguson, M., Schranz, D., 2005. Measuring default risk premia from default swap rates and edfs Unpublished working Paper, Carnegie Mellon University.

Boudoukh, J., Michaely, R., Richardson, M., Roberts, M. R., 2007. On the importance of measuring payout yield: Implications for empirical asset pricing. Journal of Finance 62, 877–915.

Breusch, T., Pagan, A., 1979. A simple test for heteroscedasticity and random coefficient variation. Econometrica 47, 1287–94.

Campbell, J. Y., 1999. Asset prices, consumption, and the business cycle. In: Taylor, J. B., Woodford, M. (eds.), *Handbook of Macroeconomics*, Elsevier, vol. 1, Part C, chap. 19, pp. 1231–1303, first ed.

Campbell, J. Y., 2003. Consumption-based asset pricing. In: Constantinides, G., Harris, M., Stulz, R. M. (eds.), *Handbook of the Economics of Finance*, Elsevier, vol. 1, Part 2, chap. 13, pp. 803–887, first ed.

Campbell, J. Y., Thompson, S. B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies 21, 1509–1531.

Caouette, J. B., Altman, E. I., Narayanan, P., 1998. Managing credit risk: the next great financial challenge, vol. 2. John Wiley & Sons.

Chen, L., Lesmond, D., Wei, J., 2007. Corporate yield spreads and bond liquidity. Journal of Finance 62, 119–149.

Chevalier, E., 2005. Critical price near maturity for an american option on a dividend-paying stock in a local volatility model. Mathematical Finance 15, 439–463.

Clarida, R., Gal, J., Gertler, M., 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. Quarterly Journal of Economics 115, 147–180.

Cochrane, J., Piazzesi, M., 2005. Bond risk premia. American Economic Review 95, 138–160.

Collin-Dufresne, P., Goldstein, R., Martin, J., 2001. The determinants of credit spread changes. Journal of Finance 56, 2177–2207.

Cooper, I., Priestley, R., 2009. Time-varying risk premiums and the output gap. Review of Financial Studies 22, 2801–2833.

Crosbie, P., Bohn, J., 2003. Modeling default risk Unpublished working paper, Moody's KMV.

David, A., 2008. Inflation uncertainty, asset valuations, and the credit spreads puzzle. Review of Financial Studies 21, 2487–2534.

Davydenko, S. A., 2012. When do firms default? a study of the default boundary Unpublished working Paper, University of Toronto.

Davydenko, S. A., Strebulaev, I. A., Zhao, X., 2012. A market-based study of the cost of default. Review of Financial Studies 25, 2959–2999.

Driessen, J., 2005. Is default event risk priced in corporate bonds? Review of Financial Studies 18, 165–195.

Duffee, G., 2007. Are variations in term premia related to the macroeconomy Unpublished working paper, Johns Hopkins University.

Duffie, D., Saita, L., Wang, K., 2007. Multi-period corporate default prediction with stochastic covariates. Journal of Financial Economics 83, 635–665.

Elton, E., Gruber, M., Agrawal, D., Mann, C., 2001. Explaining the rate spread on corporate bonds. Journal of Finance 56, 247–277.

Engelberg, J., Manski, C. F., Williams, J., 2011. Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition. Journal of Applied Econometrics 26, 1059–1078.

Eom, Y., Helwege, J., Huang, J., 2004. Structural models of corporate bond pricing: An empirical analysis. Review of Financial Studies 17, 499–544.

Epstein, L. G., Schneider, M., 2003. Recursive multiple-priors. Journal of Economic Theory 113, 1–31.

Epstein, L. G., Schneider, M., 2007. Learning under ambiguity. Review of Economic Studies 74, 1275–1303.

Fama, E., Bliss, R., 1987. The information in long-maturity forward rates. American Economic Review 77, 680–692.

Fama, E. F., Schwert, G. W., 1977. Asset returns and inflation. Journal of Financial Economics 5, 115–146.

Feldhütter, P., Schaefer, S., 2018. The myth of the credit spread puzzle. Review of Financial Studies p. forthcoming.

Fouque, J.-P., Papanicolaou, G., Sircar, K. R., 2001. From the implied volatility skew to a robust correction to black-scholes american option prices. International Journal of Theoretical and Applied Finance 4, 651–675.

Giesecke, K., Longstaff, F. A., Schaefer, S., Strebulaev, I., 2011. Corporate bond default risk: A 150-year perspective. Journal of Financial Economics 102, 233–250.

Hacker, R. S., 2010. The effectiveness of information criteria in determining unit root and trend status Unpublished working paper, Royal Institute of Technology, CESIS.

Huang, J.-Z., Huang, M., 2012. How much of the corporate-treasury yield spread is due to credit risk? Review of Asset Pricing Studies 2, 153–202.

Illeditsch, P. K., 2011. Ambiguous information, portfolio inertia, and excess volatility. Journal of Finance 66, 2213–2247.

Jones, C. S., Tuzel, S., 2013. New orders and asset prices. Review of Financial Studies 26, 115–157.

Koenker, R., 1981. A note on studentizing a test for heteroscedasticity. Journal of Econometrics 17, 107–112.

Lamberton, D., Villeneuve, S., 2003. Critical price near maturity for an american option on a dividend-paying stock. Annals of Applied Probability 13, 800–815.

Lee, R. W., 2001. Implied and local volatilities under stochastic volatility. International Journal of Theoretical and Applied Finance 4, 45–89.

Leippold, M., Trojani, F., Vanini, P., 2008. Learning and asset prices under ambiguous information. Review of Financial Studies 21, 2565–2597.

Lettau, M., Ludvigson, S., 2001. Consumption, aggregate wealth, and expected stock returns. Journal of Finance 56, 815–849.

Levendorskii, S., 2008. American and european options in multi-factor jump-diffusion models, near expiry. Finance and Stochastics 12, 541–560.

Longstaff, F. A., Schwartz, E. S., 2001. Valuing american options by simulation: A simple least-squares approach. Review of Financial studies 14, 113–147.

Ludvigson, S. C., Ng, S., 2009. Macro factors in bond risk premia. Review of Financial Studies 22, 5027–5067.

Mueller, P., Vedolin, A., Zhou, H., 2011. Short-run bond risk premia Unpublished working Paper, London School of Economics.

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Ou, S., Chlu, D., Metz, A., 2011. Corporate default and recovery rates, 1920–2010 Special Comment, Moodys Investors Service.

Piazzesi, M., Schneider, M., 2006. Equilibrium yield curves. NBER Macroeconomics Annual 2006, Volume 21 pp. 389–472.

Reisz, A. S., Perlich, C., 2007. A market-based framework for bankruptcy prediction. Journal of Financial Stability 3, 85–131.

Schaefer, S., Strebulaev, I., 2008. Structural models of credit risk are useful: Evidence from hedge ratios on corporate bonds. Journal of Financial Economics 90, 1–19.

Scholz, F., Stephens, M., 1987. Journal of the American Statistical Association 82, 918–924.

Shumway, T., 2001. Forecasting bankruptcy more accurately: A simple hazard model. Journal of Business 74, 101–24.

Sircar, R. K., Papanicolaou, G. C., 1999. Stochastic volatility, smile & asymptotics. Applied Mathematical Finance 6, 107–145.

Taylor, J. B., 1999. A Historical Analysis of Monetary Policy Rules. In: *Monetary Policy Rules*, National Bureau of Economic Research, Inc, NBER Chapters, pp. 319–348.

Vassalou, M., Xing, Y., 2004. Default risk in equity returns. Journal of Finance 59, 831–868.

Wilcox, D. W., 1992. The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work. American Economic Review 82, 922–941.

Yang, W., 2012. Business cycles and regime-shift risk Unpublished working paper, Indiana University.

Table IA.A1: Monte Carlo Simulations Based on the Learning Model: the Benchmark Specification

Panel A: Benchmark Parameterization

σ_c	$\bar{\mu}$	κ_x	σ_x	σ_y	σ_ν	ϕ_c	σ_ξ	ϕ_x
1.03	1.97	0.73	0.56	0.81	0.80	0.25	1.20	0.20

Panel B: Model-Implied Finite-Sample Properties of A_t

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Stats	Mean $\times 10^2$	Var $\times 10^4$	β $\times 10^4$	ADF Test	ACF (1)	BP Test	γ_1 $\times 10^2$	Skew-ness	VR 2-qrt	VR 10-qrt
mean	0.898	0.618	−0.003	−7.530	0.832	8.466	0.062	1.013	1.036	2.028
	(0.088)	(0.076)	(0.060)	(1.046)	(0.261)	(5.919)	(0.023)	(0.390)	(0.116)	(0.372)
p-value			0.237	0.001	0.001	0.005	0.015	0.015		
			(0.147)	(0.000)	(0.000)	(0.036)	(0.093)	(0.070)		

Panel C: Model-Implied Population Properties of A_t

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
			Kolmogorov-Smirnov tests				5-Sample	Estimates of Eq. (10)		
Stats	Mean $\times 10^2$	Var $\times 10^4$	50 vs 100	100 vs 500	500 vs 1000	1000 vs 5000	AD Test	$\hat{\kappa_a}$	$\hat{\bar{A}}$ $\times 10^2$	$\hat{\sigma_a}$ $\times 10^2$
mean	0.889	0.658	0.019	0.014	0.011	0.013	2.994	0.295	0.876	7.537
	(0.028)	(0.059)						(0.135)	(0.174)	(1.521)
p-value			0.337	0.708	0.940	0.805	0.666	0.016	0.000	0.002
								(0.007)	(0.006)	(0.008)

This table summarizes the results of Monte Carlo simulations based on the benchmark specification of the learning model, in which the agent is uncertain about the quality of public and private signals. Panel A lists the values of parameters used in simulations. Values of $\{\sigma_c, \bar{\mu}_c, \kappa_x, \sigma_x, \sigma_y\}$ are obtained from a maximum likelihood estimation of the joint dynamics of $(\log C_t, x_t, y_t)'$ as specified in Eq. (IA.A1) and (IA.A3), where changes in C_t are measured by the real GDP growth retrieved from NIPA table 1.1.1 that was last revised on December 21, 2017, x_t is measured using Blue Chip consensus forecast of one-quarter-ahead real GDP growth, and changes in y_t correspond to the second release (the one after the "preliminary" release) of real GDP growth made by the BEA. The estimation sample spans the period from 1985Q1 to 2010Q4. Panel B reports statistics of the ambiguity level $A_{x,t}$ based on 5,000 finite-sample simulations; that is, 104 quarters of data are generated for each simulation. Entries in the rows labeled "mean" and "p-value" are calculated by averaging the level and p-value of statistics across simulated samples. The corresponding cross-sample standard deviations are given in parentheses. Panel C presents results from 5,000 long-sample simulations; the length of each simulation is 5,000 quarters. Columns (4)-(7) report the results of two-sample Kolmogorov-Smirnov tests, using the simulated distributions of $A_{x,t}$ with $t = 50, 100, 500, 1000$, and 5000; Column (8) reports the corresponding results of five-sample Anderson-Darling tests. Columns (9)-(11) show the point estimates of the stochastic process in Eq. (10) using simulated samples.

Table IA.A2: Monte Carlo Simulations Based on the Learning Model: Alternative Specifications

Panel A: Model-Implied Finite-Sample Properties of A_t

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Stats	Mean $\times 10^2$	Var $\times 10^4$	β $\times 10^4$	ADF Test	ACF (1)	BP Test	γ_1 $\times 10^2$	Skew-ness	VR 2-qrt	VR 10-qrt
mean	0.952	0.662	-0.007	-7.731	0.790	6.219	0.056	0.751	1.027	1.870
	(0.076)	(0.064)	(0.060)	(0.963)	(0.227)	(6.069)	(0.026)	(0.414)	(0.107)	(0.313)
p-value			0.234	0.001	0.001	0.008	0.023	0.048		
			(0.146)	(0.000)	(0.030)	(0.093)	(0.086)	(0.087)		

Panel B: Model-Implied Population Properties of A_t

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
			Kolmogorov-Smirnov tests				5-Sample	Estimates of Eq. (10)		
Stats	Mean $\times 10^2$	Var $\times 10^4$	50 vs 100	100 vs 500	500 vs 1000	1000 vs 5000	AD Test	$\hat{\kappa_a}$	$\hat{\bar{A}}$ $\times 10^2$	$\hat{\sigma_a}$ $\times 10^2$
mean	0.977	0.684	0.015	0.009	0.012	0.009	2.522	0.306	0.967	7.668
	(0.024)	(0.052)						(0.090)	(0.144)	(1.095)
p-value			0.641	0.987	0.888	0.984	0.769	0.039	0.000	0.001
								(0.013)	(0.010)	(0.007)

Panel C: Model Implications for the Uncertainty Level in Contemporaneous Growth Rates

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Finite-Sample				Population				Data			
Mean $\times 10^2$	Var of $\times 10^4$	Skew-ness	ACF (1)	Mean $\times 10^2$	Var of $\times 10^4$	Skew-ness	ACF (1)	Mean $\times 10^2$	Var of $\times 10^4$	Skew-ness	ACF (1)
1.058	0.642	1.065	0.615	1.059	0.628	1.089	0.579	1.088	0.657	1.228	0.635
(0.031)	(0.109)	(0.409)	(0.168)	(0.014)	(0.037)	(0.224)	(0.261)				

This table summarizes the results of Monte Carlo simulations based on an alternative specification of the learning model: the agent has learned the precision of the public signal but is still uncertain about the quality of private signals; the latter are specific to priors. Panel A reports statistics of the ambiguity level $A_{x,t}$ based on 5,000 finite-sample simulations; that is, 104 quarters of data are generated for each simulation. Entries in the rows labeled "mean" and "p-value" are calculated by averaging the level and p-value of statistics across simulated samples. The corresponding cross-sample standard deviations are given in parentheses. Panel B presents results from 5,000 long-sample simulations; the length of each simulation is 5,000 quarters. Columns (4)-(7) report the results of two-sample Kolmogorov-Smirnov tests, using the simulated distributions of $A_{x,t}$ with $t = 50, 100, 500, 1000$, and 5000; Column (8) reports the corresponding results of five-sample Anderson-Darling tests. Columns (9)-(11) show the point estimates of the stochastic process in Eq. (10) using simulated samples. Panel C shows the properties of the model-implied dispersion in contemporaneous projections of realized growth rates as well as their empirical counterparts; the former are derived from 5,000 simulations and the latter are from the interdecile range of forecasters' projections for the current quarter.

Table IA.B3: Survey Data Summary and Description

Yr-Qtr	No_BC	Pct_BC	No_SPF	Pct_SPF	Yr-Qtr	No_BC	Pct_BC	No_SPF	Pct_SPF	Yr-Qtr	No_BC	Pct_BC	No_SPF	Pct_SPF
1985Q1	45		24	75%	1993Q4	47	94%	34	79%	2002Q3	49	96%	35	91%
1985Q2	48	94%	29	69%	1994Q1	50	94%	30	83%	2002Q4	48	98%	35	83%
1985Q3	47	100%	31	71%	1994Q2	47	96%	35	80%	2003Q1	47	96%	37	78%
1985Q4	48	98%	25	84%	1994Q3	47	96%	30	87%	2003Q2	46	96%	35	86%
1986Q1	46	98%	26	77%	1994Q4	50	92%	30	73%	2003Q3	47	96%	30	80%
1986Q2	46	98%	28	75%	1995Q1	49	98%	29	83%	2003Q4	47	94%	34	71%
1986Q3	43	100%	24	71%	1995Q2	48	98%	52	52%	2004Q1	46	100%	32	81%
1986Q4	45	91%	24	71%	1995Q3	48	96%	52	85%	2004Q2	46	96%	32	78%
1987Q1	46	94%	22	86%	1995Q4	49	88%	47	87%	2004Q3	46	98%	30	87%
1987Q2	50	92%	28	71%	1996Q1	49	92%	38	92%	2004Q4	44	95%	34	79%
1987Q3	50	100%	22	91%	1996Q2	47	94%	42	74%	2005Q1	47	91%	36	86%
1987Q4	44	100%	20	80%	1996Q3	49	92%	42	79%	2005Q2	48	94%	46	59%
1988Q1	48	96%	18	72%	1996Q4	49	98%	37	84%	2005Q3	50	94%	53	79%
1988Q2	47	96%	19	74%	1997Q1	49	98%	38	82%	2005Q4	50	98%	51	94%
1988Q3	49	94%	18	78%	1997Q2	47	96%	40	83%	2006Q1	50	96%	53	87%
1988Q4	48	98%	15	80%	1997Q3	47	96%	37	86%	2006Q2	49	92%	53	92%
1989Q1	49	96%	17	71%	1997Q4	48	92%	38	82%	2006Q3	50	94%	51	86%
1989Q2	49	98%	17	71%	1998Q1	45	98%	32	94%	2006Q4	50	94%	51	86%
1989Q3	47	100%	15	73%	1998Q2	48	91%	29	76%	2007Q1	51	98%	49	86%
1989Q4	50	94%	17	71%	1998Q3	48	88%	32	69%	2007Q2	51	100%	53	81%
1990Q1	48	100%	14	64%	1998Q4	45	96%	33	73%	2007Q3	50	100%	49	88%
1990Q2	48	98%	9	78%	1999Q1	47	91%	33	76%	2007Q4	48	100%	48	81%
1990Q3	47	98%	13	54%	1999Q2	46	91%	37	65%	2008Q1	47	96%	50	84%
1990Q4	47	96%	30	33%	1999Q3	48	90%	37	81%	2008Q2	46	96%	50	84%
1991Q1	48	94%	35	63%	1999Q4	46	89%	43	70%	2008Q3	46	98%	47	83%
1991Q2	48	96%	42	79%	2000Q1	47	91%	36	83%	2008Q4	45	98%	51	76%
1991Q3	50	96%	33	91%	2000Q2	44	91%	34	76%	2009Q1	46	87%	43	91%
1991Q4	48	100%	39	79%	2000Q3	46	93%	32	84%	2009Q2	46	89%	51	78%
1992Q1	49	96%	40	90%	2000Q4	46	98%	34	82%	2009Q3	48	90%	34	94%
1992Q2	48	98%	37	95%	2001Q1	43	98%	35	83%	2009Q4	48	94%	41	66%
1992Q3	48	96%	36	83%	2001Q2	46	89%	33	94%	2010Q1	49	96%	42	90%
1992Q4	48	98%	40	83%	2001Q3	47	100%	33	82%	2010Q2	48	100%	44	89%
1993Q1	49	−−*	34	94%	2001Q4	42	95%	29	79%	2010Q3	46	100%	36	92%
1993Q2	50	94%	36	78%	2002Q1	49	86%	35	63%	2010Q4	48	94%	43	74%
1993Q3	49	98%	35	83%	2002Q2	47	96%	37	84%					

* *In January 1993, the Blue Chips changed the coding system used to identify individual forecasters.*

This table provides descriptive results on the panelists participating in the Blue Chip Financial Forecasts and the Survey of Professional Forecasters. The columns labeled "No_BC" and "No_SPF" report the number of respondents contained in the survey of each period. The columns labeled "Pct_BC" and "Pct_SPF" show the percentage of current panel members that also responded to the last period's survey. The sample period spans from 1985Q1 to 2010Q4.

Table IA.B4: Empirical Properties of the Ambiguity Level

Panel A: Variance Ratios Tests					Panel B: Properties of Skewness	
Horizon(yr)	VR(i)	90%	95%			
2	1.232	1.108	1.142		Sample	0.820
5	1.841	1.191	1.278		90%	0.504
10	2.363	1.241	1.398		95%	0.358

Panel A shows the variance ratios (VR) for the absolute value of the residuals from a AR(5) model for the ambiguity measure,

$$\widetilde{A}_t = \sum_{j=1}^{5} B_j L^j \widetilde{A}_t + \epsilon_t,$$

where L is a lag operator. Bootstrap percentiles are computed by resampling the residual series 10,000 times, under the null hypothesis that realized volatility were i.i.d. Panel B presents skewness estimates and corresponding bootstrap confidence percentiles. The sample period spans from 1985Q1 to 2010Q4.

Table IA.B5: Additional Predictors of Asset Returns and Credit Spreads

	In-sample										Out-of-sample	
	Ambi-guity	Net Payout	Detrended $y_t^{(0.25)}$	cay	$\Delta y_t^{(10)}$	S&P$_t$	ΔVIX_t	Cochrane-Piazzesi	Ludvigson-Ng	Adjusted R^2	Out-of-Sample R^2	Clark-West p-value
Panel A: Predictability in excess stock market returns												
r_{t+1}^M	2.945	0.035								0.036	0.029	0.019
	(2.924)	(0.932)										
r_{t+1}^M	3.977		-0.607							0.036	0.027	0.006
	(3.090)		(-1.127)									
r_{t+1}^M	2.928			0.636						0.037	0.042	0.038
	(2.568)			(1.368)								
Panel B: Predictability in corporate yield spreads												
CS_{t+1}^{IG}	0.409				-0.210					0.122	0.081	0.030
	(2.530)				(-1.717)							
CS_{t+1}^{IG}	0.410					-0.050				0.213	0.082	0.039
	(2.564)					(-1.614)						
CS_{t+1}^{IG}	0.649						0.025			0.192	0.179	0.001
	(2.031)						(1.225)					
CS_{t+1}^{HY}	2.966				-1.502					0.351	0.372	0.001
	(3.212)				(-2.265)							
CS_{t+1}^{HY}	2.646					-0.182				0.357	0.303	0.000
	(2.759)					(-2.367)						
CS_{t+1}^{HY}	2.992						0.086			0.326	0.359	0.000
	(3.099)						(1.312)					
Panel C: Predictability in excess Treasury bond returns												
r_{t+1}^{Int}	0.992							0.166		0.074	0.022	0.026
	(2.588)							(2.943)				
r_{t+1}^{Int}	1.069								0.101	0.078	0.045	0.047
	(2.385)								(3.029)			
r_{t+1}^{long}	1.318							0.373		0.066	0.024	0.027
	(2.052)							(3.205)				
r_{t+1}^{long}	1.778								0.207	0.058	0.062	0.029
	(1.938)								(2.998)			

This table contains results from regressing (one-quarter-ahead) excess stock returns, credit spreads, and excess bond returns on the proposed ambiguity measure and other predictors. Conditioning predictors of equity returns include net payout yields, detrended risk-free rates (measured as three-month T-bill rates), and Lettau and Ludvigson (2001)'s *cay*. Alternative predictors of credit spreads include changes in risk-free rates (measured as 10-year T-note rates), S&P 500 returns, and changes in the VIX index. Alternative predictors of Treasury bond returns include Cochrane and Piazzesi (2005)'s return-forecasting factors as well as Ludvigson and Ng (2009)'s macroeconomic factor. r_t^M is defined as the difference between the continuously compounded return on the CRSP value-weighted index and the contemporaneous return on a three-month Treasury bill. CS_t^{IG} and CS_t^{HY} denote the credit spreads of Barclays US aggregate corporate and high-yield bond indices. r_{t+1}^{Int} and r_{t+1}^{Long} are quarterly returns on the Barclays long-term and intermediate-term Treasury indices in excess of the contemporaneous return on a three-month Treasury bill. t-statistics are computed following the procedures of Newey and West (1987) with a lag truncation parameter of 6. The last two columns report Campbell and Thompson (2008)'s out-of-sample R^2s as well as p-values of the Clark-West (2007) test. The benchmark model in the Clark-West test is based on the historical average return (spread) or one of conditioning predictors. The sample period spans from 1985Q1 to 2010Q4.

Table IA.C6: On the Estimation of the Cost of Default and Default Boundary

Panel A: Number of Defaults by Year

	Defaulted bonds	Defaulting firms	Defaulting public firms
1986	25	7	1
1987	26	11	1
1988	43	15	5
1989	46	12	2
1990	87	42	9
1991	119	49	9
1992	41	17	6
1993	17	13	1
1994	13	10	2
1995	26	20	4
1997	24	15	2
1998	34	24	3
1999	105	65	20
2000	100	63	12
2001	195	84	19
2002	103	47	14
2003	102	42	5
2006	20	10	2
2007	7	3	4
2008	116	55	45
2009	214	82	43
2010	16	10	10
All	1479	696	219

Panel B: Estimates of the Default Cost and Default Boundary

	Mean	Median	Std	10%	90%
Asset return (%)	−15.9	−16.4	15.0	−34.9	2.0
Equity return (%)	−33.6	−26.2	29.4	−77.6	−2.1
Debt return (%)	−13.4	−10.6	16.8	−40.1	5.4
Debt recovery rate (%)	45.6	43.4	23.9	15.5	79.2
Default cost (%)	24.7	29.3	25.9	−2.0	58.4
Default boundary (%)	67.6	69.6	28.4	28.0	104.8

This table summarizes the sample and results in our estimation of the cost of default and solvancy boundary. Panel A lists the number of defaulted bond issues and defaulting issuers in our sample by year. Panel B shows various statistics describing how the firm's value changes in the month of default. The reported asset returns are adjusted by subtracting the return on the S&P 500 over the same month. The cost of default is expressed as a fraction of the market value of unlevered equity at the end of the last calendar month prior to default, and the default boundary a fraction of the face value of debt.

Table IA.E7: Empirical Credit Spreads of Ten-Year Bonds

	Aaa	Aa	A	Baa	Ba	B
1985-2010:						
Individual Bonds	77	91	117	182	298	488
Barclays Idx.	95	112	147	210	390	580
1985-1997:						
Individual Bonds	70	87	106	153	276	459
Barclays Idx.	84	109	125	179	340	532
1998-2010:						
Individual Bonds	80	95	131	198	334	506
Barclays Idx.	106	115	169	242	420	609
BofA ML Idx.	74	100	132	194	372	

This table shows the levels of historical 10-year credit spreads across ratings categories. For rows named "Individual Bonds", average yield spread is firstly calculated for each bond with the maturity between 7 and 15 years, and the median value is then taken across bonds. Individual bond's quotes are obtained from Lehman Brothers fixed-income data set for the 1985-1997 period, and from Merrill Lynch for the 1998-2010 period. The "Barclays Idx." rows report the average credit spreads of Barclays U.S. Aggregate Corporate/High Yield Bond Indices, and the "BofA ML Idx." rows show the means of Bank of America Merrill Lynch US corporate/high-yield 7-10Y option-adjusted spreads.

Table IA.E8: Mean and Median Spreads in the Merton Model

		Panel A: Mean-to-Mean			Panel B: Median-to-Median		
		1y	4y	10y	1y	4y	10y
AAA	Model_CS	0	8	12	0	0	0
	Model_Rep	{0}	{0}	{4}	{0}	{0}	{1}
	Actual	⟨32⟩	⟨68⟩	⟨90⟩	⟨28⟩	⟨64⟩	⟨77⟩
AA		11	30	31	0	7	12
		{0}	{6}	{18}	{0}	{5}	{16}
		⟨41⟩	⟨85⟩	⟨104⟩	⟨41⟩	⟨83⟩	⟨91⟩
A		16	35	44	0	5	25
		{1}	{10}	{27}	{1}	{9}	{25}
		⟨54⟩	⟨102⟩	⟨134⟩	⟨53⟩	⟨99⟩	⟨117⟩
Baa		57	91	99	3	54	59
		{4}	{36}	{65}	{2}	{47}	{56}
		⟨113⟩	⟨170⟩	⟨207⟩	⟨91⟩	⟨139⟩	⟨182⟩
Ba		100	144	193	17	104	134
		{23}	{151}	{182}	{16}	{101}	{127}
		⟨248⟩	⟨306⟩	⟨357⟩	⟨195⟩	⟨274⟩	⟨298⟩
B		279	353	376	183	234	296
		{236}	{344}	{390}	{179}	{227}	{303}
		⟨391⟩	⟨480⟩	⟨539⟩	⟨317⟩	⟨424⟩	⟨488⟩

This table examines convexity bias when credit spreads in the Merton model is derived from a representative firm. "Model_CS" denotes the average/median model spreads over the cross section of individual bonds: the Merton spread is firstly computed for each bond-month observation; the time series for a specific bond is then divided by the credit rating, and a simple average is then taken over the monthly spreads in the same rating category; finally the means/medians of these averaged spreads are calculated within each rating group. Model spreads in angle brackets are based on the "representative firm" calculation: for each rating category we calculate the average/median leverage ratio, asset volatility as well as payout ratio, and the model spread is then computed using the average/median firm variables. The mean/median of observed spreads are reported in braces. Individual bond's prices are obtained from Lehman Brothers fixed-income data set for the 1985-1997 period, and from Merrill Lynch for the 1998-2010 period. For an industrial bond to enter the sample, it must belong to a Lehman or Merrill Lynch index and must have at least 24 monthly trader quotes over the 1985-2010 period.

Fig. IA.B1. **Density Functions of the Amount of Ambiguity** A_t

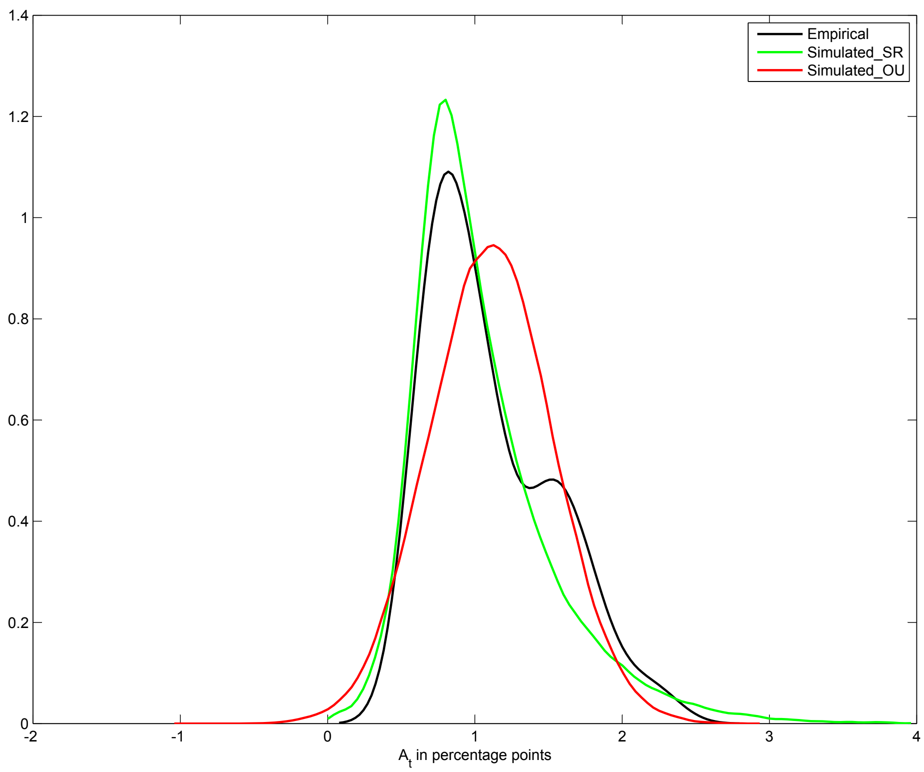

This figure shows empirical densities of A_t (black lines) and simulated densities (green and red lines). "Simulated_SR" denotes the simulated density of the estimated mean-reverting square-root process, $dA_t = \kappa(\bar{A} - A_t)dt + \sigma_a\sqrt{A_t}dB_{A,t}$, and "Simulated_OU" the corresponding Ornstein-Uhlenbeck process with $dA_t = \kappa(\bar{A} - A_t)dt + \sigma_a dB_{A,t}$.