Long-Run Risks, Credit Markets, and Financial Structure

By Harjoat S. Bhamra, Lars-Alexander Kuehn, and Ilya A. Strebulaev*

In this paper we explore the impact of long-run risks in cash flow and consumption growth on optimal corporate default and capital structure decisions, the term structure of credit spreads and actual default probabilities, and the levered equity risk premium. We do this by embedding a structural model of credit risk and dynamic corporate financing decisions in a consumption-based representative-agent asset pricing model. The resulting unified framework has the advantage that it can be used to study the interplay between corporate finance and asset pricing. In contrast with many previous studies, which consider an individual firm at its refinancing point, our framework incorporates firm-level heterogeneity in capital structure decisions via a dynamic cross-section of firms. Since empirically firms are known to change their financial structure infrequently, modeling a dynamic cross-section is crucial for understanding how financing decisions affect asset prices.

To study long-run risks, we tailor approach by assuming the representative agent has Epstein-Zin-Weil preferences and that the first and second conditional moments of earnings and consumption growth are slowly mean-reverting. We find that long-run risk lowers firms' optimal leverage ratios, while also increasing both credit spreads and the levered equity risk premium. The basic intuition is that the Epstein-Zin-Weil agent prefers uncertainty about the first and second moments of growth rates to be resolved sooner rather than later, and so, the probability of expected growth rates being low and growth rate volatilities being high is greater in the risk-neutral compared to the actual world. Hence, risk-neutral default rates are high, while actual rates are low. Also, the agent prices assets as if bad states last longer than is actually the case, which raises the risk premium.

Rather than restricting our analysis to credit spreads and actual default probabilities for one maturity, we consider the term structure. It is here that firm-level heterogeneity is of first order importance. An individual firm's term structure of default probabilities is much steeper than in the historical data and the five-year default probability is particularly small. When we compute average default probabilities for a cross-section of firms, two changes occur. First, the resulting term structure is higher at five years and much closer to the data. Second, the term structure is less steep, and so the ten-year actual default probability is also much closer to the data. The first effect stems from positive skewness in the long-run distribution of firm cash flows, which increases the number of firms close to default. The second effect depends on the evolution of the distribution of firm cash flows. Over time, firms' earnings on average increase which reduces positive skewness, diminishing the mass of firms close to the default boundary. Hence, the difference between ten- and five-year default probabilities is smaller in the cross-section than for an individual firm.

Our calibration assumes that expected macroeconomic growth rates and volatilities are highly persistent. Thus, the risks inherent in changing macroeconomic growth rates are long-run risks as in Ravi Bansal and Amir Yaron (2004). This is in contrast with Harjoat S. Bhamra, Lars-A. Kuehn and Ilya A. Strebulaev (2009), and Harjoat S. Bhamra, Lars-A. Kuehn and Ilya A. Strebulaev (Forthcoming), where expected macroeconomic growth rates and volatilities revert to their means at rates that corre-

* Bhamra: Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver BC, Canada V6T 1Z2, harjoat.bhamra@sauder.ubc.ca. Kuehn: Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, kuehn@cmu.edu. Strebulaev: Graduate School of Business, Stanford University, 518 Memorial Way, Stanford, CA 94305, istrebulaev@stanford.edu. We thank Ravi Bansal, Adlai Fisher, Stephen Schaefer, Raman Uppal, Toni Whited, and Amir Yaron for helpful comments and discussion.

spond more closely to business cycles. Empirical evidence on expected macroeconomic growth rates suggests they contain both business cycle and long-run risk components (e.g., Stephen G. Cecchetti, Pok-sang Lam and Nelson C. Mark (1993) for evidence on business cycle frequency components). This suggests that it may be fruitful to use a multi-frequency approach (see Laurent E. Calvet and Adlai J. Fisher (2001)) to jointly model the asset pricing and credit risk implications of these two distinct components.

I. Model

Aggregate consumption, C_t, is given by

$$\frac{dC_t}{C_t} = g_t dt + \sigma_{C,t} dB_{C,t}, \tag{1}$$

where g_t is expected consumption growth, $\sigma_{C,t}$ is consumption growth volatility, and $B_{C,t}$ is a Brownian motion. The are N firms in the economy. The earnings process for firm i, $X_{i,t}$ is given by

$$\frac{dX_{i,t}}{X_{i,t}} = \theta_{i,t} dt + \sigma^{id}_{X,i} dB^{id}_{X,i,t} + \sigma^{s}_{X,i,t} dB^{s}_{X,t}, \tag{2}$$

where θ_i is the expected earnings growth rate of firm i, and $\sigma^{id}_{X,i}$ and $\sigma^{s}_{X,i}$ are, respectively, the idiosyncratic and systematic volatilities of the firm's earnings growth rate. Total risk, $\sigma_{X,i}$, is given by $\sigma_{X,i} = \sqrt{(\sigma^{id}_{X,i})^2 + (\sigma^{s}_{X,i})^2}$. Although firms account correctly for total risk in making financial decisions, the idiosyncratic component generates cross-sectional heterogeneity in risks over time. The Brownian motion $B^{s}_{X,t}$ is the systematic shock to the firm's earnings growth, which is correlated with aggregate consumption growth:

$$dB^{s}_{X,t} dB_{C,t} = \rho_{XC} dt, \tag{3}$$

where ρ_{XC} is the correlation coefficient. The Brownian motion $B^{id}_{X,i,t}$ is the idiosyncratic shock to firm i's earnings, which is correlated with neither $B^{s}_{X,t}$, $B_{C,t}$, nor with other firms' idiosyncratic shocks.

We assume there is long-run cash flow risk, as in Bansal and Yaron (2004), by assuming that g_t, θ_t, $\sigma_{C,t}$, and $\sigma^{s}_{X,t}$ depend on the state of the economy, ν_t, which can take 2 values.[1] We assume state 1 is the bad state of the economy. Since the first (second) moments of fundamental growth rates are procyclical (countercyclical), $g_1 < g_2$, $\theta_1 < \theta_2$, $\sigma_{C,1} > \sigma_{C,2}$, and $\sigma^{s}_{X,1} > \sigma^{s}_{X,2}$.

The state changes according to a 2-state Markov chain, defined by λ_{ν_t}, $\nu_t \in \{1, 2\}$, which is the probability per unit time of the economy leaving state ν_t. The Markov chain gives rise to uncertainty about the future moments of consumption growth. This intertemporal consumption risk impacts the state-price density only if the representative agent cares about the intertemporal distribution of risk. To ensure this, we assume the representative agent has the continuous-time analog of Epstein-Zin-Weil preferences (see Larry G. Epstein and Stanley E. Zin (1989)). Consequently, the representative agent's state-price density at time-t, π_t, is given by

$$\pi_t = \left(\beta e^{-\beta t}\right)^{\frac{1-\gamma}{1-\frac{1}{\psi}}} C_t^{-\gamma} \left(p_{C,t} e^{\int_0^t p_{C,s}^{-1} ds}\right)^{-\frac{\gamma-\frac{1}{\psi}}{1-\frac{1}{\psi}}}, \tag{4}$$

where β is the rate of time preference, γ is the coefficient of relative risk aversion (RRA), and ψ is the elasticity of intertemporal substitution under certainty (EIS).[2] The Epstein-Zin-Weil agent cares whether news about consumption growth and hence future consumption is good or bad. Her state-price density then, unlike that of the power-utility agent, depends on the value of the claim to aggregate consumption per unit consumption, i.e. the price-consumption ratio, p_C. We assume $\gamma > 1/\psi$, which implies the agent prefers intertemporal risk to be resolved sooner rather than later. Consequently, bad news about consumption growth increases the state-price density, as can be seen from (4).[3]

[1] To ensure idiosyncratic earnings volatility, σ^{id}_X, is truly idiosyncratic, we assume it is constant and thus independent of the state of the economy. We also assume that the correlation coefficient, ρ_{XC}, is constant.

[2] The continuous-time version of recursive preferences is known as stochastic differential utility (see Darrell Duffie and Larry G. Epstein (1992)).

[3] Later, to ensure p_C is procyclical, we shall impose the additional assumption that $\psi > 1$.

II. Dynamic Capital Structure

We follow standard EBIT-based capital structure models (e.g., Robert Goldstein, Nengjiu Ju and Hayne Leland (2001)) and assume the earnings of a firm, X, are split between a coupon, c, promised to debt holders in perpetuity and a dividend, $X - c$, paid to equityholders. The after-tax distribution to equityholders is therefore $(1-\eta)(X-c)$ where η denotes the tax rate. Equityholders of each individual firm make three types of corporate financing decisions: (i) they have the right to default at the time of their choice; (ii) they decide when to refinance the firm's debt obligations; and (iii) they decide on the amount of debt to be issued at each refinancing.

As is well-known, equityholders exercise their default option if earnings drop below a certain earnings level, called the default boundary. The default boundary in our framework is endogenously state-dependent. Specifically, default occurs when a firm's cash flow level reaches a lower boundary, $X_{D,\nu_0\nu_D}$, where ν_0 is the state at the most recent refinancing date and ν_D is the state of the economy at default. Upon default, bondholders recover a proportion of the firm's assets in lieu of coupons, i.e., a fraction α_{ν_t} of the after-tax present value of the firm's earnings.

At date 0, equityholders choose how much debt to issue, by selecting the coupon, c_{ν_0}, which depends on the state of the economy at that date. At later times, equityholders can choose to restructure their existing debt obligations. In common with the literature we assume that refinancings are leverage increasing transactions since empirical evidence demonstrates that reducing leverage in distress is much more costly. Firms also prefer to refinance infrequently since each refinancing is costly Edwin O. Fischer, Robert Heinkel and Josef Zechner (1989).

Refinancing occurs when earnings reach an upper boundary, $X_{U,\nu_0\nu_U}$, which again depends on ν_0 and the state of the economy at refinancing, ν_U. At each refinancing, equityholders choose a new coupon to maximize their value. For simplicity, we assume that debt is non-callable and issued *pari passu*, i.e., all outstanding debt issues have equal seniority. Dilution is on a per-coupon basis, so that if the coupon at the previous refinancing is c_{ν_0}, and the new coupon is $c_{\nu_U}(c_{\nu_0})$, then the old debtholders are still entitled to the coupon, c_{ν_0}, and their stake in future default and refinancings is $c_{\nu_0}/c_{\nu_U}(c_{\nu_0})$. This introduces path-dependence of capital structure in the sense that historical macroeconomic conditions affect current capital structure.

III. Results

Although most of the results can be obtained in closed form, we report here the results of the calibration to highlight the empirical relevance of the model. We choose parameter values so that the statistical properties of consumption and earnings growth match, as closely as possible, those reported in Ravi Bansal, Dana Kiku and Amir Yaron (2007). We calibrate idiosyncratic earnings volatility so that the total asset volatility is approximately 25 percent, the average asset volatility of firms with outstanding rated corporate debt. This yields an idiosyncratic earnings volatility of 22.84 percent per annum. As in Bhamra, Kuehn and Strebulaev (2009) and Bhamra, Kuehn and Strebulaev (Forthcoming), we assume bankruptcy costs, $1 - \alpha_{\nu_t}$, are countercyclical. See Table 1 for details.

Table 1—Calibration

	State 1	State 2
Consumption growth	0.0071	0.0289
Consumption volatility	0.0260	0.0048
Earnings growth	-0.0071	0.0431
Systematic earnings volatility	0.1718	0.0315
Idiosyncratic earnings volatility	0.2284	0.2284
Earnings/consumption growth correlation	0.5755	0.5755
Long-run probabilities	0.5000	0.5000
Convergence to long-run	0.2153	0.2153
Annual discount rate	0.01	0.01
Tax rate	0.15	0.15
Bankruptcy costs	0.30	0.10
Debt issuance cost	0.03	0.01
Relative risk aversion	10	10
EIS	2	2

First, we investigate the credit risk puzzle. The key aspects of our methodology are that we (i) compute model-implied credit spreads as the *average* credit spread over the *cross-section* of firms, (ii) pay particular attention to the impact of the *time evolution* of the distribution on the *term structure* of actual default probabilities and credit spreads. Our results are consistent with empirically observed credit spreads and actual default probabilities at several maturities, as shown in Table 2. Panel A shows results for an individual firm at the refinancing point for the specification with an optimal default boundary and exogenous leverage. Panel B reports the results for a dynamic simulated cross-section of BBB firms with optimal default and dynamic capital structure decisions. The default rate is the cumulative default rate over the horizon of five or ten years. Moments of the default rate distribution are based on 1,000 generated economies. Credit spreads are given in basis points and all other variables in per cent.

Table 2—Credit risk implications

	Panel A:		Panel B:	
Maturity	5	10	5	10
Credit spread	16	85	120	150
Default rate				
Average	0.1	1.2	1.0	2.8
Median			0.7	2.4
25% quantile			0.2	1.6
75% quantile			1.6	4.2
Leverage	40	40	Sample	Sample

As our results show, for an *individual* firm, actual default probabilities are too low and their term structure is too steep relative to the data. For the *distribution* of firms, average actual default probabilities are higher and the term structure flatter, thus matching the data more closely. The reason for the higher average actual default probabilities based on the distribution of firms is that some firms will be near default and actual default probabilities are convex in the distance-to-default. The intuition for the flattening of the term structure is quite different and is based on the *time evolution* of the distribution of firms. Specifically, the right tail of the distribution of firms' cash flows (and hence distances-to-default) is fatter at ten years than five years, and thus the slope of the term structure of average default probabilities for a distribution of firms is flatter than the term structure of default probabilities for an individual firm.

Next, we study the implications of the model for the levered equity premium. First, we investigate the impact of default and refinancing options, when leverage is endogenous. Second, we study the impact of aggregation. Panel A of Table 3 shows how leverage, with the default and refinancing options present, affects the levered equity risk premium at refinancing, whereas Panel B shows the effects of aggregation.

Table 3—The levered equity risk premium

In Panel A, the cases are: (i) no leverage; (ii) leverage without default and refinancing options; (iii) leverage with the default option but without the refinancing option; (iv) for the benchmark case of leverage with both default and refinancing options.

Panel A				
	(i)	(ii)	(iii)	(iv)
Risk Prem.	3.51	11.82	4.23	4.40
Leverage	0.00	66.22	27.65	23.20

Panel B		
	Re Date	True Dyn
Equity premium (un)	3.51	3.51
Equity premium (le)	4.40	4.53
Leverage	23.20	40.44

Panel A shows that the levered equity risk premium is always higher than the unlevered equity risk premium (3.51 percent). In particular, the risk premium in the absence of default and refinancing options is 11.82 percent (compared to 8.10 percent in Bansal, Kiku and Yaron (2007)). However, we also find that the default option lowers leverage and hence *reduces* the levered equity risk premium to 4.23 percent. The intuition is that costly default makes shareholders issue less debt and thus lowers the probability of default. Similarly, introducing the refi-

nancing option lowers leverage, since equityholders now have a real option of adding more debt later. Therefore, this option reduces the levered risk premium, giving a benchmark value of 4.40 percent.

In Panel B we compute the risk premium at the refinancing date under dynamic capital structure (Re Date), and in the cross-section under dynamic capital structure, i.e. 'true dynamics' (True Dyn). There are two dimensions to our results: the impact of leverage and the impact of a dynamic cross-section of firms, i.e. aggregation. Leverage has two effects on the risk premium. First, the dividend payment to equityholders becomes riskier, which raises the risk premium. Second, the presence of a default option (driven by limited liability) shifts value from debtholders to equityholders, decreasing the premium. Panel B shows that the first effect dominates and leverage increases the risk premium from 3.51 percent to 4.40 percent.

Panel B shows that the levered equity risk premium under 'true dynamics' is 4.53 percent. Thus, accounting for aggregation raises the levered equity premium. The reason is that in the dynamic cross-section, most firms are not at their refinancing points, and for these firms leverage in general is not equal to leverage at the most recent refinancing point. In fact, average leverage is higher. The intuition is quite general: unsuccessful firms refinance later than successful firms and, as a result, have higher leverage ratios, especially so because firms which opt for higher leverage at refinancing also choose a lower refinancing boundary. Consequently, the risk premium under 'true dynamics' is higher than at the refinancing point.

The equity premium obtained under the long-run risk calibration is higher than under the business cycle calibration used in Bhamra, Kuehn and Strebulaev (2009). With the long-run risk calibration, macroeconomic growth rates are more persistent, leading to greater intertemporal risk, and hence a higher premium.

IV. Conclusion

We investigate the impact of long-run risks in cash flow and consumption growth on credit risk and the levered equity risk premium by building a unified framework that embeds a dynamic capital structure model in a consumption-based asset pricing model. Our results show that the model can produce both a reasonable equity risk premium and term structure of credit spreads, while keeping default rates realistically low.

REFERENCES

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59(4): 1481–1509.

Bansal, Ravi, Dana Kiku, and Amir Yaron. 2007. "Risks For the Long Run: Estimation and Inference." http://faculty.fuqua.duke.edu/ rb7/bio.

Bhamra, Harjoat S., Lars-A. Kuehn, and Ilya A. Strebulaev. 2009. "The Levered Equity Risk Premium and Credits Spreads: A Unified Approach." *Review of Financial Studies.*

Bhamra, Harjoat S., Lars-A. Kuehn, and Ilya A. Strebulaev. Forthcoming. "The Aggregate Dynamics of Capital Structure and Macroeconomic Risk." *Review of Financial Studies.*

Calvet, Laurent E., and Adlai J. Fisher. 2001. "Forecasting Multifractal Volatility." *Journal of Econometrics*, 105: 27–58.

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark. 1993. "The Equity Premium and the Risk-Free Rate: Matching the Moments." *Journal of Monetary Economics*, 31(1): 21–45.

Duffie, Darrell, and Larry G. Epstein. 1992. "Stochastic Differential Utility." *Econometrica*, 60(2): 353–94.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4): 937–69.

Fischer, Edwin O., Robert Heinkel, and Josef Zechner. 1989. "Dynamic Capital Structure Choice: Theory and Tests." *Journal of Finance*, 44(1): 19–40.

Goldstein, Robert, Nengjiu Ju, and Hayne Leland. 2001. "An EBIT-Based Model of Dynamic Capital Structure." *Journal of Business*, 74(4): 483–512.