Ambiguity Shifts and the 2007–2008 Financial Crisis☆

Nina Boyarchenko*

Federal Reserve Bank of New York

Abstract

Faced with doubts about the quality of information and the quality of modeling techniques, ambiguity-averse agents assign higher probabilities to lower utility states, leading to higher CDS premia and lower equity prices. Using data on financial institutions, I find that the sudden increases in credit spreads during the recent crisis can be explained by changes in the amount of ambiguity faced by market participants and changes in how the total amount of ambiguity was distributed between ambiguity about information quality and ambiguity about model quality.

JEL codes: G01, G12, G13

Keywords: Ambiguity aversion, credit risk, financial crisis

☆This paper is based on my dissertation completed at the University of Chicago, Booth School of Business and Department of Economics. I would like to thank my committee Lars P. Hansen (co-chair), Pietro Veronesi (co-chair), Stavros Panageas and John Heaton. I would also like to thank Lubos Pastor, Martin Schneider, Valentin Haddad, Matthew Plosser, John Crosby, Paul Whelan (EFA discussant), Itamar Drechsler (Carnegie-Rochester discussant), Hui Chen (AFA discussant), seminar participants at University of Chicago (Booth School of Business), Johns Hopkins University (Carey School of Business), Federal Reserve Board of Governors, Federal Reserve Bank of New York, Ohio State University (Fisher College of Business), the Einaudi Institute for Economics and Finance and Bocconi University and participants in EFA Stockholm, Carnegie-Rochester Conference November 2011, and AFA Chicago for helpful comments on the previous drafts of this paper. All the remaining errors are mine. Research support from the Sanford J. Grossman Fellowship in Honor of Arnold Zellner is gratefully acknowledged. Any opinions expressed herein are the author's and are not representative of the views of Sanford J. Grossman, Arnold Zellner, the Federal Reserve Bank of New York or of the Federal Reserve System.

*Corresponding author, 33 Liberty Street, New York, NY 10045, email: nina.boyarchenko@ny.frb.org, phone: (212)-720-7339

1. Introduction

When making consumption decisions, an investor faces uncertainty about both the relevant underlying state and the data-generating process governing the evolution of the state. While uncertainty about the state is risk that the investor understands and can model, uncertainty about the data-generating process represents agents' pessimism about their ability to identify the correct model. This paper argues that prices of credit securities are sensitive to the investors' preferences toward model uncertainty and that the implied time-variation in the level of model uncertainty is a source of variation in credit spreads that explains the asymmetric response of credit spreads to upturns and downturns in the economy.

The recent financial crisis highlighted two short-comings of standard structural credit risk models. First, models calibrated to match credit spreads at long maturities underestimate credit spreads at short maturities. Second, rapid changes in the overall level of credit spreads cannot be explained in the standard risk aversion framework.[1] Matching the levels of short-term credit spreads necessitates that the structural model can generate non-zero default probabilities even at short maturities while matching the overall time series evolution of the credit curve requires that the model can update default probabilities quickly. One of the difficulties in explaining the behavior of credit spreads during the recent episode is the speed with which market conditions deteriorated. At the beginning of 2007, financial markets were liquidity-unconstrained and credit spreads were at historical lows. Even as late as May 2007, it would have been hard to predict the magnitude of the response that the losses on subprime mortgages had generated. Compared to the total value of financial instruments traded worldwide, the subprime losses were relatively small: even the worst-case estimates put them at around USD 250 billion.[2] Further, for investors familiar with these instruments, the losses were not unexpected. By definition, the subprime mortgages were part of the riskiest segment of the mortgage market, so it was hardly surprising some borrowers would default on the loans. Yet, despite their predictability, the defaults had precipitated the current liquidity crisis that spread between the credit markets. This effect is particularly clear in the evolution of the five year default swap (CDS) spread on major financial institutions during the crisis, plotted in Fig. 1. Before the start of the crisis in August 2007, the five year CDS rate for these companies was below 50 bps. During the course of the crisis, it increased to as high as 1000 bps.

[1] See Online Appendix for an example of calculating default swap spreads with risk averse agents.
[2] Source: Caballero and Krishnamurthy (2008)

In this paper, I argue that the increases in CDS spreads observed during the 2007–2008 crisis were due to increases in investors' doubts about the validity of their pricing models and the quality of the signals available to market participants. I modify the standard Black and Cox (1976) model of credit risk by introducing ambiguity-averse, imperfectly-informed agents. The ambiguity-aversion reflects the fact that, when it is difficult for investors to observe firms' assets directly, they are forced to rely on imprecise accounting information. In this situation, investors must draw inference from accounting data and other publicly available information. Investors realize that, although they may be able to pick a model of the fundamental asset value and the accounting signals to best fit the historical data, this may not be the true data-generating model. Under this assumption, I derive the asset prices in the economy, explicitly accounting for the implications of imperfect information and model misspecification.

The ambiguity aversion preferences representative agent impact credit spreads in two ways. First, since ambiguity-averse agents place a higher probability of transitioning to a lower utility state next period, they believe the probability of default next period to be higher than their ambiguity-neutral counterparts, raising the short term credit spreads relative to the benchmark model. In the model considered in this paper, this effect is further exacerbated by the presence of imperfect information. As shown in Duffie and Lando (2001), imperfect information also raises the level of short term credit spreads since imperfectly-informed agents are uncertain about how close a firm is to its default boundary. Ambiguity-aversion has a bite here as well: ambiguity-averse agents interpret signals pessimistically and believe the firm to be closer to the default boundary to begin with, raising the short-term credit spreads even further relative to the perfectly-informed, ambiguity-averse agents.

Since ambiguity-averse agents assign higher probabilities to lower utility outcomes, they also react more strongly to negative news than agents who use Bayes' Law to update their beliefs. In particular, on observing unexpected negative signals, ambiguity-averse agents can quickly update their beliefs about default probabilities, leading to rapid increases in CDS spreads. Notice that, when agents are averse to ambiguity about both signal quality and model quality, agents can use negative signals to either update their beliefs about the current state or about the underlying transition dynamics. Thus, the strong reaction to negative signals does not necessarily constitute a panic on the part of investors.

Using observations of the CDS spreads on financial institutions, I estimate the degrees of investors' aversion to ambiguity about signal quality and to ambiguity about the fundamental asset value dynamics. Estimating the two ambiguity aversion coefficients using three sub-periods – before the start of the crisis in July 2007,

from the start of the crisis to the bailout of Bear Stearns in March 2008, and from the bailout of Bear Stearns to the liquidation of Lehman Brothers in September 2008 – I find that the three estimates are not statistically significantly different, implying that investors' misspecification preferences did not change during the crisis. Instead, changes in the level of CDS spreads during the crisis were driven by increases in the amount of ambiguity in the economy. Furthermore, the relative contribution of the two sources of ambiguity changed during the crisis. The initial BNP Paribas announcement in August 2007 and the subsequent Northern Rock revelations lead to an increase in ambiguity about the quality of the signals available to market participants. Intuitively, while no major financial institution had defaulted on its obligations, market participants could interpret the signals in the market as informative about the quality of the signals produced by financial institutions, and not about the quality of the assets on their balance sheets. The bailout of Bear Stearns and the bailout of Lehman Brothers, on the other hand, lead to an increase in ambiguity about the underlying dynamics. That is, the effective default of these two institutions increased investors' doubts about the quality of the surviving financial institutions. The level of ambiguity stabilizes after the introduction of Troubled Asset Relief Program (TARP) in October 2008 but at a higher level than before the start of the crisis.

Using the estimates from the CDS market, I then compute the model-implied equity prices for the financial institutions. While the credit risk model cannot match the levels of equity prices and, in fact, is not geared to do so, it does match the qualitative behavior of equity prices during the crisis. Since equity is a long-lived claim to a firm's cash-flows, changes in the amount of ambiguity about signal quality have a smaller impact on equity than on CDS spreads. Thus, the model with ambiguity averse agents can generate the non-decrease in equity prices of financial institutions before the bailout of Bear Stearns in March 2008 that we observe in the data. Notice that this fact is a challenge even for models with time-varying risk aversion as the increase in risk aversion necessary to explain the increases in CDS spreads would be accompanied by decreases in equity prices.

The rest of the paper is organized as follows. I review the related literature in Section 2. I describe the model considered in the paper in Section 3. The results of the estimation of the model are presented in Section 4. Section 5 concludes. Technical details are relegated to the appendix.

2. Literature Review

A rapidly growing literature studies the behavior of asset prices in the presence of ambiguity in dynamic economies. A substantial part of this literature considers

investor ambiguity about the data-generating model. Anderson et al. (2003) derive the pricing semigroups associated with robust perturbations of the true state probability law. Trojani and Vanini (2002) use their framework to address the equity premium and the interest rate puzzles, while Leippold et al. (2008) consider also the excess volatility puzzle. Gagliardini et al. (2009) study the term structure implications of adding ambiguity to a production economy. This setting has also been used to study the portfolio behavior of ambiguity-averse investors and the implications for the options markets (see e.g. Trojani and Vanini (2004), Liu et al. (2005), Drechsler (2010)).

The second strand in the literature, however, assumes that, although the agents in the economy know the "true" data-generating model, they face uncertainty about the quality of the observed signal about an unobservable underlying. Chen and Epstein (2002) study the equity premium and the interest rate puzzles in this set-up, and Epstein and Schneider (2008) consider the implications for the excess volatility puzzle. The portfolio allocation implications of this setting have also been studied extensively in e.g. Uppal and Wang (2003) and Epstein and Miao (2003).

However, none of these papers study the relationship between ambiguity aversion and the term structure of credit spreads. Following Hansen and Sargent (2005, 2007), I introduce model misspecification by considering martingale distortions to the reference model probability law. As Hansen and Sargent (2007) show, the martingale distortion can be factored into distortions of the conditional distribution of the underlying state (signal quality) and the evolution law of the hidden state (asset value dynamics). I assume that the representative investor in the secondary debt market has max-min preferences over consumption paths under possible models.

This paper is also related to the literature on preference-based explanations for the credit spread puzzle (that yield spreads between corporate bonds and treasuries are high and volatile). One popular approach is to examine credit spreads in the presence of long-run risk. Chen (2010) uses a discrete-state version of Bansal and Yaron (2004) model of long-run risk as a driver of business cycles. The paper argues that the credit spread puzzle and the under-leverage puzzle can be explained by two observations: defaults are more highly concentrated during economy downturns, when marginal utility is high, and the losses associated with default are higher during such times. Thus, investors demand high risk premia for holding defaultable claims, including corporate bonds and levered firms. Bhamra et al. (2010) also use long-run risk as a driver of business cycles and show that the optimal earnings default boundary, default rates and credit spreads are countercyclical.

Using similar intuition, Chen et al. (2009) argue that the credit spread puzzle can be explained by the covariation between default rates and market Sharpe ratios.

That is, investors must be compensated more for holding credit risk securities because default rates (and, hence, expected losses from default) increase at the same time as market returns are more uncertain. More specifically, the authors investigate the credit spread implications of the Campbell and Cochrane (1999) pricing kernel calibrated to equity returns and aggregate consumption data. Identifying the historical surplus–consumption ratio from aggregate consumption data, the paper finds that the implied level and time-variation of spreads match historical levels well.

3. Model

In this section, I present the economy considered in this paper. I begin by describing the reference model used for pricing credit securities and then proceed to the misspecification problem faced by the representative agent in the economy.

3.1. Reference model

As the reference model, I consider a modified version of the Black and Cox (1976) economy. Consider a (sector of the) economy consisting of I firms, indexed by $i = 1, \ldots, I$ and denote by $A_{it} = e^{a_{it}}$ the fundamental value of the assets of firm i at date $t = 1, 2, \ldots$. To fix ideas, assume that there are $n_y = 12$ data periods in a year, so that each period corresponds to a month. I assume that the log-asset value of each firm can be decomposed into the sum of two components:

$$a_{it} = z_{it} + \rho_i z_{ct}, \tag{1}$$

where z_{it} is a firm-specific shock to the asset value of firm i, z_{ct} is an aggregate shock to the asset values of all the firms in the sector and ρ_i is the loading of firm i on the aggregate component. Denote by $z_t = [z_{1t}, \ldots, z_{It}, z_{ct}]'$ the vector of the components of asset values at date t. I assume that the vector z_t evolves according to an N-state Markov chain, with possible values $\xi_1, \ldots, \xi_N$ and the transition probability matrix Λ defined as:

$$\{\Lambda\}_{jk} \equiv \lambda_{jk} = \mathbb{P}\left(z_{t+1} = \xi_k \middle| z_t = \xi_j \right). \tag{2}$$

There are two types of agents in the economy: managers and investors. All the day-to-day operations of the firm are delegated to the respective manager. I assume that there are no agency problems between a firm's managers and the equity holders of the firm, so that the managers act in the best interest of the equity holders. Further, similarly to Duffie and Lando (2001), I assume that the managers are better informed about the firm they manage than the participants in the public markets and, in particular, that the managers of the firm observe perfectly the evolution

of the fundamental value of the firm's assets. To prevent information spill-over, I assume that managers are precluded from trading in the public assets markets.

In this paper, I abstract from modeling the operational decisions of the firm managers and, in particular, from modeling the optimal capital structure, dividend payment and default decisions faced by the managers. As in Leland (1994) and Duffie and Lando (2001), I assume that each firm i issues perpetual debt with face value D_i. This debt is serviced by a constant coupon rate C_i. While firm i is in operation, it generates a constant fraction δ_i of assets as cash-flows which accrue, minus the coupon payments, as equity in the firm.

Notice that the above specification excludes any time trend in firms' earnings. An alternative would be to specify the time t value of firm i as $A_{it} = e^{\gamma_i t + a_{it}}$, where γ_i is the deterministic trend for firm i. The stochastic component $e^{a_{it}}$ then represents deviation in firm value from the deterministic trend. With this specification, the face value of debt would also grow with the deterministic trend, so that the constant D_i would translate into a constant leverage ratio.

The managers decide on behalf of the equity holders when to default. As in Black and Cox (1976), I abstract from modeling the liquidation decision faced by the managers and assume instead that, for each firm i, there is a state i^* of the Markov chain that is labeled as the default state for that firm. Denoting by $a_{B_i} = \xi_{i^*}$ the vector of asset values in state i^*, the (stochastic) default date τ_i of firm i is the first hitting time of the state a_{B_i}: $\tau_i = \inf\{t : z_t = a_{B_i}\}$. In economic terms, the exogenous default rule can be interpreted as a debt covenant. The firm is liquidated at the present value of the discontinued cash flows, with the proceeds distributed among the firm's primary debt holders and the equity holders receiving 0. Although this assumption implies that there is no cost of default, it is possible to incorporate liquidation costs into the model by allowing the value for firm i in its default state to reflect a cost of liquidation. For simplicity, I assume that each firm has a single default state and that firms do not default simultaneously. Notice that, since managers observe perfectly the asset value evolution of the firm under their management, there is no uncertainty about the firm being liquidated upon hitting its default boundary. Finally, denote by $a_B = \bigcup_{i=1}^{I} a_{B_i}$ the union of the default states of all the firms and by a_B^c its complement, which is the set of states where none of the firms default.

Consider now the participants in the public markets. Similarly to Duffie and Lando (2001), I assume that the representative investor does not observe the true evolution of asset values in the sector and receives instead imperfect, unbiased signals about the fundamental value of the assets of each firm, $\hat{A}_{it} = e^{y_{it}}$, and the aggregate component of asset values in the sector, $\hat{A}_{ct} = e^{y_{ct}}$. More specifically, assume that $y_{it} = a_{it} + u_{it}$ and $y_{ct} = z_{ct} + u_{ct}$ where the signal errors $u_t = [u_{1t}, \ldots, u_{It}, u_{ct}]'$ are

serially uncorrelated and normally distributed, independent of the true realization of z_t: $u_t \sim N(\overline{u}, \Sigma_u)$. Here, $\overline{u}$ is the mean signal error and Σ_u^{-1} the signal quality. At each date t, the representative agent also observes whether any of the default states have been reached and any of the firms have been liquidated. Thus, the information set of the representative agent at date t is:

$$\mathcal{G}_t = \sigma\left\{y_s,\ \mathbf{1}_{z_s \in a_B^c}:\ s = 1, \ldots, t\right\},$$

where $y_t = [y_{1t}, \ldots, y_{It}, y_{ct}]'$ is the full signal vector at date t.

Denote by p_{jt} the probability, conditional on the date t information set of the representative investor, of the vector z being in state j at date t:

$$p_{jt} = \mathbb{P}\left(z_t = \xi_j \middle| \mathcal{G}_t\right).$$

In the Online Appendix, I present the updating rule in terms for the vector of unnormalized probabilities $\vec{\pi}_t$, which are related to the proper probabilities by:

$$p_{jt} = \frac{\pi_{jt}}{\sum_{j=1}^{N} \pi_{jt}}, \qquad j = 1, \ldots, N.$$

Finally, consider the utility of the representative agent. I assume the representative agent is risk-neutral and, thus, holds all the claims to the firm's asset value. Thus, the date t expected present value of the utility of the representative agent is given by:

$$J_t = \mathbb{E}\left[\sum_{s=0}^{+\infty} \beta^s \sum_{i=1}^{I} \delta_i A_{i,t+s} \middle| \mathcal{G}_t\right], \tag{3}$$

where β is the subjective discount factor. For the discussion below, it is useful to represent the expected present value of utility in recursive form:

$$J_t = \mathbb{E}\left[\sum_{i=1}^{I} \delta_i A_{it} + \beta \mathbf{1}_{z_{t+1} \in a_B^c} J_{t+1} + \beta \sum_{i=1}^{I} \mathbf{1}_{z_{t+1} = a_{B_i}} J^{\tau_i} \middle| \mathcal{G}_t\right], \tag{4}$$

where J^{τ_i} is the value function of the representative agent in case firm i defaults.

3.2. Asset prices

In this paper, I consider two types of claims to the assets of the firm: a claim to the firm's equity and a credit default swap (CDS) written on the consol bond. Recall that equity in firm i accrues as a constant fraction δ_i of the fundamental asset value

and that the equity holders receive 0 in case of the firm being liquidated. Thus, the date t price of a claim to equity of firm i, V_{it}, is given by:

$$V_{it} = \mathbb{E}\left[\sum_{s=0}^{\tau_i} \beta^s(\delta_i A_{i,t+s} - C_i)\middle| \mathcal{G}_t\right],$$

where τ_i is the (stochastic) default date of the firm i. Notice that the equity price satisfies the Euler equation:

$$V_{it} = \mathbb{E}\left[\delta_i A_{it} - C_i + \beta \mathbf{1}_{\tau_i > t+1} V_{i,t+1} \middle| \mathcal{G}_t\right]. \tag{5}$$

Consider now the default swaps written on the primary debt of firm i. With a given maturity T, a default swap is an exchange of an annuity stream at a constant coupon rate until maturity or default, whichever is first, in return for a payment of X at default, if default is before T, where X is the difference between the face value and the recovery value on the stipulated underlying bond. A default swap can thus be thought of as a default insurance contract for bond holders that expires at a given date T, and makes up the difference between face and recovery values on the bond in the event of default.

I assume, as typical in practice, that the default swap annuity payments are made semiannually, and that the default swap's maturity date T is a coupon date. As in Duffie and Lando (2001), I take the underlying bond for the default swap on firm i to be the consol bond issued by firm i. Recall that, in case of default, the debt holders receive the present value of the discontinued cash flows. Thus, the payment X_i per unit of primary debt if firm i defaults before the swap maturity date T is given by:

$$X_i = 1 - \frac{\delta_i (I - \beta\Lambda)^{-1}_{i^*} e^{\xi}}{D_i}.$$

That is, the payment in case of default compensates the protection buyers for income lost due to the default of the underlying entity. The at-market default swap spread is the annualized coupon rate $c_i(t,T)$ that makes the default swap sell at date t for a market value of 0. Thus, with $T = t + 6n$ for a given positive integer n,[3] the CDS spread is given by:

$$c_i(t,T) = \frac{2X_i \mathbb{E}\left[\beta^{\tau_i - t}\mathbf{1}_{\tau_i < T} \middle| \mathcal{G}_t\right]}{\sum_{s=1}^{n} \beta^{6s} \mathbb{E}\left[\mathbf{1}_{\tau_i \geq t+6s} \middle| \mathcal{G}_t\right]}. \tag{6}$$

[3]Recall that there are 12 data periods in a year

In the above, the numerator reflects the expected present value of payment to the insurance buyer in case of default before the maturity date T, while the denominator reflects the present value of the scheduled payments to the insurance seller until default or maturity, whichever one comes first.

3.3. Model misspecification

This paper studies asset prices in a setting where the representative agent makes decision rules robust to possible misspecifications of asset value and accounting signals models. In reality, the correct specification assumption of the reference model is overly restrictive. It implies that, even though the participants in the public markets only observe imperfect signals about the evolution of the fundamental asset value, they can still correctly identify the parametric model that governs the relevant dynamics. More realistically, I assume that the representative investor in the firm fears misspecification of the probability law generated by the model above and believes instead that the signals are related to the true asset value realizations by a family of likelihoods.

As in Hansen and Sargent (1995), Hansen et al. (1999) and Anderson et al. (2003), I model preferences of the representative agent in the presence of model misspecification using the recursion:

$$J_t = -\theta_s \log \mathbb{E}\left[\exp\left[-\frac{U(z_t) + \mathcal{R}_t(\beta J_{t+1};\theta_d)}{\theta_s}\right]\middle|\, \mathcal{G}_t\right], \tag{7}$$

where:

$$\mathcal{R}_t(\beta J_{t+1};\theta_d) \equiv -\theta_d \log \mathbb{E}\left[\exp\left(-\frac{\beta J_{t+1}}{\theta_d}\right)\middle|\, \mathcal{F}_t\right]$$

and

$$\mathcal{F}_t = \sigma\left\{z_s,\ y_s,\ \mathbf{1}_{z_s \in a_B^c} :\ s = 1, \ldots, t\right\}$$

is the full information set. The tilted recursion (7) replaces the standard utility recursion (4), incorporating the representative agent's misspecification doubts in two steps. First, the tilted continuation function $\mathcal{R}_t$ makes an additional risk adjustment to the continuation value function of the representative agent, accounting for misspecification fears about the fundamental asset value evolution dynamics. Second, the tilted expectations over the current period utility adjusts for misspecification fears about the filtered probability distribution over the current state. As emphasized by Hansen and Sargent (1995), the log-exp specification of the recursion links risk-sensitive control theory and a more general recursive utility specification of Epstein and Zin (1989). The degree of the representative agent's aversion to mis-

specification of the underlying dynamics is quantified by θ_d^{-1} (where d stands for dynamics) and the degree of the representative agent's aversion to misspecification of the filter distribution is quantified by θ_s^{-1} (where s stands for signals). When $\theta_s^{-1} = \theta_d^{-1} = 0$, the risk-sensitive recursion (7) reverts to the usual utility recursion under Von Neumann-Morgenstern form of state additivity. For values of θ_d^{-1} greater than zero, the recursion (7) implies an increased aversion to risk associated with the time evolution of the hidden state *vis a vis* the Von Neumann-Morgenstern specification. Similarly, for values of θ_s^{-1} greater than zero, the recursion implies an increased aversion to risk associated with the unobservability of the true state *vis a vis* the Von Neumann-Morgenstern specification. Finally, notice that, since the distorted continuation value function $\mathcal{R}_t$ conditions on the full information set $\mathcal{F}_t$ (and, hence, on the true realization of the hidden state), the investors in this economy have the option of focusing their attention on misspecification of the joint probability distribution of future signals and state.

Hansen and Sargent (2007) show that, corresponding to the tilted continuation function $\mathcal{R}_t$ is the worst-case likelihood ratio:

$$\phi_t(z_{t+1}, y_{t+1}) = \frac{\exp\left(-\frac{\beta J_{t+1}}{\theta_d}\right)}{\mathbb{E}\left[\exp\left(-\frac{\beta J_{t+1}}{\theta_d}\right)\middle|\, \mathcal{F}_t\right]}. \tag{8}$$

ϕ_t tilts the joint distribution toward lower continuation value states, decreasing the expected future value of the continuation utility. Similarly, corresponding to the recursion (7), is the worst-case likelihood ratio:

$$\psi_t(z) = \frac{\exp\left(-\frac{U(z)+\mathcal{R}_t(\beta J_{t+1};\theta_d)}{\theta_s}\right)}{\mathbb{E}\left[\exp\left(-\frac{U(z)+\mathcal{R}_t(\beta J_{t+1};\theta_d)}{\theta_s}\right)\middle|\, \mathcal{G}_t\right]} \tag{9}$$

between the conditional distribution over the state at date t under the misspecfied and reference models. ψ_t tilts the conditional distribution toward lower utility states, decreasing the expected value of utility at the current date.

To quantify the amount of distortion in the economy, Hansen and Sargent (2007) introduce measures of the conditional relative entropy between the reference and the distorted models. In particular, under the full information setting, the conditional relative entropy between the reference and the worst case model of the future state

and signals evolution is defined as:

$$\epsilon_t^1(\phi_{t+1}, \xi_j) = \sum_{k=1}^{N} \int \tau(\xi_k, y_{t+1}|\xi_j)\phi(\xi_k, y_{t+1}|\xi_j) \log \phi(\xi_k, y_{t+1}|\xi_j) dy_{t+1}$$

and the conditional relative entropy between the reference and the worst case model of the current state by:

$$\epsilon_t^2(\psi_t) = \sum_{j=1}^{N} p_j \psi_j \log \psi_j.$$

The total conditional relative entropy between the reference model and the worst case model at date t is then given by:

$$\epsilon_t = \epsilon_t^2(\psi_t) + \tilde{\mathbb{E}}\left[\epsilon_t^1(\phi_{t+1}, z_t)\middle|\, \mathcal{G}_t\right] \equiv \epsilon_t^2(\psi_t) + \hat{\epsilon}_t^1(\phi_{t+1}). \tag{10}$$

In general, the recursion (7) does not have a closed-form solution. Instead, I look for a first order approximation to the representative agent's value function around the point $\theta_d^{-1} = \theta_s^{-1} = 0$, which corresponds to the solution under the reference model. Notice that the approximation I construct here is different in its nature from the small noise approximations constructed in Campi and James (1996) and Anderson et al. (2012) as, instead of approximating the reference model value function around the deterministic steady state, I approximate around the value function corresponding to the zero signal precision case.[4] The following result holds:

Lemma 3.1. *The first order approximation to the value function around the point $\theta_d^{-1} = \theta_s^{-1} = 0$ is given by;*

$$J(\pi; \theta_d^{-1}, \theta_s^{-1}) \approx J_0(\pi) + \theta_d^{-1} J_{\theta_d^{-1}}(\pi) + \epsilon J_\epsilon(\pi), \tag{11}$$

where $\epsilon = \theta_s^{-1}/\theta_d^{-1} - 1$. The first order approximation to the implied distortion to the conditional joint distribution of next period's signals and state is then given by:

$$\begin{aligned} \phi_t\left(z_{t+1}, y_{t+1}\right) \approx 1 &- \theta_d^{-1}\beta\left(J_0(\pi_{t+1}) - \mathbb{E}\left[J_0(\pi_{t+1})\middle|\, \mathcal{F}_t\right]\right) \\ &- \left(\theta_s^{-1} - \theta_d^{-1}\right)\beta\left(J_\epsilon(\pi_{t+1}) - \mathbb{E}\left[J_\epsilon(\pi_{t+1})\middle|\, \mathcal{F}_t\right]\right), \end{aligned} \tag{12}$$

and the first order approximation to the implied distortion to the conditional distri-

[4]When the signal precision approaches 0, the agent does not update the conditional probability distribution and, hence, the stationary distribution can be used as the prior distribution.

bution of the current state by:

$$\psi_t(z_t) \approx 1 - \left(\theta_d^{-1} - \theta_s^{-1}\right) \beta \left(\mathbb{E}\left[J_\epsilon(\pi_{t+1}) \middle| \mathcal{F}_t\right] - \mathbb{E}\left[J_\epsilon(\pi_{t+1}) \middle| \mathcal{G}_t\right]\right) \\ - \theta_s^{-1} \left(U(z_t) + \beta \mathbb{E}\left[J_0(\pi_{t+1}) \middle| \mathcal{F}_t\right] - \mathbb{E}\left[U(z_t) + \beta J_0(\pi_{t+1}) \middle| \mathcal{G}_t\right]\right). \tag{13}$$

Proof. See Online Appendix. □

Here, J_0 is the value function of the ambiguity-neutral agent, $J_{\theta_d^{-1}}$ measures the change in the value function as the overall level of ambiguity in the economy increases and J_ϵ measures the change in the value function as the agent becomes more averse to ambiguity about the signal quality than ambiguity about the underlying dynamics. Since the ambiguity-averse agent is solving a minimization problem in determining optimal consumption and portfolio choices, $J_{\theta_d^{-1}}$ is negative. Thus, as the agent becomes more ambiguity averse overall, his value function decreases. Similarly, J_ϵ is also negative, since increasing the agent's aversion to ambiguity about the signal quality increases the total level of ambiguity. Notice that, since both $J_{\theta_d^{-1}}$ and J_ϵ are time-varying, the agent's perceived risk attitudes change depending on the conditional distribution of the hidden state.

Consider now the implied amount of ambiguity, as measured by the relative entropy between the reference and the worst-case models. Substituting (12) into the definition of ϵ_t^1, the relative entropy between the reference and the worst-case models for the fundamental asset dynamics, we obtain:

$$\begin{aligned} \epsilon_t^1 &= \mathbb{E}\left[\phi_t(z_{t+1}, y_{t+1}) \log \phi_t(z_{t+1}, y_{t+1}) \middle| \mathcal{F}_t\right] \\ &\approx \theta_d^{-2} \beta^2 var\left(J_0\left(\pi_{t+1}\right) \middle| \mathcal{F}_t\right) + \left(\theta_s^{-1} - \theta_d^{-1}\right)^2 \beta^2 var\left(J_\epsilon\left(\pi_{t+1}\right) \middle| \mathcal{F}_t\right) \\ &+ 2\left(\theta_s^{-1} - \theta_d^{-1}\right) \theta_d^{-1} \beta^2 cov\left(J_0\left(\pi_{t+1}\right), J_\epsilon\left(\pi_{t+1}\right) \middle| \mathcal{F}_t\right). \end{aligned}$$

Notice first that, when the representative agent is equally averse to both sources of ambiguity, so that $\theta_d = \theta_s$, the amount of ambiguity due to misspecification of the fundamental asset value dynamics depends only on the variance of the next period's value function. For the case $\theta_d \neq \theta_s$, two more quantities are instrumental in determining the amount of ambiguity: the variance of next period's correction to the value function in the direction of greater aversion to ambiguity about signal quality and the covariance between the reference model value function and the correction term. Since J_0 and J_ϵ are negatively correlated, there is a trade-off between aversion to ambiguity about signal quality and aversion to ambiguity about the underlying dynamics. In particular, if the representative agent is more averse to ambiguity about signal quality than to ambiguity about the fundamental asset value dynamics

(so that $\theta_d^{-1} < \theta_s^{-1}$), aversion to ambiguity about signal quality reduces the amount of ambiguity about the underlying dynamics. If, on the other hand, the agent is more averse to ambiguity about the underlying dynamics (so that $\theta_d^{-1} > \theta_s^{-1}$), the aversion to ambiguity about signal quality increases the amount of ambiguity about the fundamental asset value. Intuitively, if the agent is more averse to ambiguity about signal quality, then he will not apply as a large of a distortion to the transition probabilities since misspecification of the underlying dynamics is a smaller concern in his mind.

Similarly, we can represent the relative entropy between the reference and the worst-case models for the posterior distribution of the current state as:

$$\begin{aligned}\epsilon_t^2 &= \mathbb{E}\left[\psi_t(z_t)\log\psi_t(z_t)\middle|\,\mathcal{G}_t\right]\\ &\approx \theta_s^{-2}var\left(U(z_t)+\beta J_0(\pi_{t+1})\middle|\,\mathcal{G}_t\right)+\left(\theta_d^{-1}-\theta_s^{-1}\right)^2\beta^2 var\left(J_\epsilon\left(\pi_{t+1}\right)\middle|\,\mathcal{G}_t\right)\\ &+2\left(\theta_d^{-1}-\theta_s^{-1}\right)\theta_s^{-1}\beta cov\left(J_\epsilon\left(\pi_{t+1}\right),U(z_t)+\beta J_0(\pi_{t+1})\middle|\,\mathcal{G}_t\right).\end{aligned}$$

Thus, when the agent is equally averse to both sources of ambiguity, the amount of ambiguity about signal quality depends only on the variance of the current period reference model value function. For the case $\theta_s \neq \theta_d$, however, the trade-off between aversion to ambiguity about signal quality and aversion to ambiguity about the fundamental asset value dynamics reemerges. If the agent is more averse to ambiguity about the signal quality than about the underlying dynamics (so that $\theta_d^{-1} < \theta_s^{-1}$), the aversion to ambiguity about the underlying dynamics increases the amount of ambiguity about signal quality. If, on the other hand, the agent is more averse to ambiguity about the underlying dynamics (so that $\theta_d^{-1} > \theta_s^{-1}$), the aversion to ambiguity about the underlying dynamics mitigates the impact of the ambiguity about the signal quality. Intuitively, if the representative agent is more averse to ambiguity about the underlying dynamics, he will not apply as large a distortion to the current period posterior probabilities since misspecification of the filter dynamics is a smaller concern in his mind.

Finally, notice that the implied transition probability matrix can be represented as:

$$\begin{aligned}\tilde{\lambda}_{jk} &\equiv \tilde{\mathbb{P}}\left(z_{t+1}=\xi_k\middle|\,z_t=\xi_j\right)\\ &\approx \lambda_{jk}\left[1-\theta_d^{-1}\left(\mathbb{E}\left[J_0(\pi_{t+1})\middle|\,\mathcal{F}_t,z_{t+1}=\xi_k\right]-\mathbb{E}\left[J_0(\pi_{t+1})\middle|\,\mathcal{F}_t\right]\right)\right]\\ &-\lambda_{jk}\left(\theta_s^{-1}-\theta_d^{-1}\right)\left(\mathbb{E}\left[J_\epsilon(\pi_{t+1})\middle|\,\mathcal{F}_t,z_{t+1}=\xi_k\right]-\mathbb{E}\left[J_\epsilon(\pi_{t+1})\middle|\,\mathcal{F}_t\right]\right).\end{aligned} \tag{14}$$

Thus, the worst-case transition probability matrix puts a higher transition proba-

bility on states with lower expected utility payoffs. Notice further that the magnitude of the distortion is determined primarily by the degree of the agent's aversion to ambiguity about the underlying dynamics, θ_d^{-1}, since the value function of the ambiguity-neutral agent dominates the correction terms. Notice also that the above implies that, unlike the reference model, the transition probability matrix under the misspecified model is time-dependent introducing additional time variation in asset prices in the economy.

3.4. Asset prices under the misspecified model

Since the representative agent evaluates expectations under the worst-case measure when making consumption decisions, the Euler equation holds under the worst-case likelihood and assets can be priced using the worst-case Euler equation. That is, the date t price of a claim to the equity of firm i is given by (5) and the corresponding CDS spread by (6) but with the expectations taken under the worst-case likelihood.

To understand the contributions of the two sources of ambiguity to asset prices, compare the price of a claim to a firm's equity and the 6 month CDS spread on the same firm. By definition, the 6 month CDS spread reflects only the probability that the firm defaults before six months. At such short horizons, the current state has a large impact in determining whether the firm will default before the specified horizon. Thus, the 6 month CDS spread will reflect the representative agent's aversion to both ambiguity about signal quality (since that influences his beliefs about the current state) and ambiguity about the underlying dynamics (since that influences how fast the economy can transition to the default state of the firm). The equity price, on the other hand, reflects the agents' beliefs about the probability of default at all future horizons. As the horizon increases, the corresponding default probability is determined more by the structure of the transition probability matrix and less by the current state. Since the short default horizons correspond to a small portion of the expected cash-flows from the equity instrument, the price of equity is determined mainly by aversion of the representative agent to ambiguity about the underlying dynamics. The rest of the CDS term structure lies between these two extremes, with the agent's aversion to ambiguity about the underlying dynamics playing an increasingly important role as the maturity of the CDS contract increases.

4. The 2007–2008 Financial Crisis

In this section, I estimate the model in Section 3 using data on financial institutions. Although the returns on the equity of financial institutions accounts for a small portion of the overall level of consumption in the economy, these institutions

were at the forefront of the 2007 financial crisis and, thus, to understand the asset price movements during the crisis, it is important to understand the movements in the prices of claims on these institutions. I begin by estimating the parameters of the reference model using the observations of book equity of financial institutions as firm-specific signals and the Case-Shiller 10 Cities Housing (CS10) Index as the signal about the common component of the asset values. Although book equity of financial institutions makes for imperfect signals, both because by its nature the series is backward-looking and because of the infrequency of observations, using observations of market equity is also fraught with difficulty. The problem of inverting market prices to obtain time series observations of risk-neutral probabilities is a complex problem in general. For the model considered in this paper, the matter is complicated by the fact that all market prices reflect the worst-case distortions imposed by the preferences of the representative agent. Thus, if market prices were to be used as firm-specific signals, then estimation could no longer be decomposed into two parts – estimating independently the parameters of the reference model and then using these parameters as inputs to estimate the ambiguity preference parameters, further complicating the estimation procedure. While I believe that observations of market equity for the institutions concerned provide interesting insights into the nature of risk and uncertainty in the market, in the current paper I prefer to use these time series to provide an outside test for the validity of the estimated parameters. I choose observations of the CS10 Index as the aggregate signal to capture the exposure of the financial institutions to risks associated with the national housing market.

Next, I use historical observations of CDS spreads for the financial institutions prior to the start of the crisis to obtain an estimate of θ_s and θ_d – the parameters governing the investors' aversion to misspecification of the filter distribution and underlying dynamics, respectively. Using these estimates, I compute the implied relative entropy between the reference and the worst-case models and decompose the entropy calculation into the contributions from misspecification of the signal model and misspecification of the fundamental value of assets model for the whole time series. Finally, I compare the model-implied CDS rates and equity prices to the observed time series.

4.1. Estimating the reference model

To estimate the reference model, I use historical observations of the firm-specific signals and the aggregate signal. The sample includes monthly observations of the Case-Shiller 10 Cities Index from May 1987 through July 2007, and quarterly observations of the balance sheet variables for the same time period. Below, I provide the outline of the estimation procedure. The details of the estimation are provided in

Online Appendix.

As observations of firm-specific signals, I use observations of book equity from COMPUSTAT. As observations of the aggregate signals, I take the time series of the Case-Shiller 10 Cities index. Notice that, while balance sheet data are observed at a quarterly frequency only, observations of the CS10 Index are available at a monthly frequency. The procedure described in Online Appendix accounts explicitly for this dual frequency of observations. Notice also that, to estimate the reference model, I only use observations up to Q2 2007 to avoid introducing measurement error by including observations of the accounting signal which reflect mark-downs taken since the start of the crisis, as well as the increased ambiguity discount in credit derivatives held on the balance sheets of these institutions.

Recall that the reference model is described by the parameters:

$$\left(\Lambda,\ \{\xi_j\}_{j=1}^{N},\ \{\rho_i\}_{i=1}^{I},\ \Sigma_u,\ \{D_i, C_i\}_{i=1}^{I},\ \{\delta_i\}_{i=1}^{I}\right).$$

I begin the estimation by identifying the face value of the perpetual bond issued by firm i, D_i, with the last pre-crisis (Q2 2007) observation of the firm's value of long-term debt; notice that, since the model-implied debt has infinite maturity, long-term debt is a better measure than total debt as it excludes short-term liabilities. The coupon payment, C_i, is then chosen to make the level of debt D_i optimal.[5] Following the model assumption that each firm generates a constant fraction δ_i of assets as cash-flows, I identify δ_i as the time-series average of the total earnings as a fraction of total assets. The rest of the parameters are estimated using the Gibbs sampling procedure of Online Appendix. The estimated reference model parameters are also presented in Online Appendix.

4.2. Estimating the misspecification preference parameter

Consider now estimating the degree of aversion toward misspecification of the underlying dynamics, θ_d^{-1}, and misspecification of the filter distribution, θ_s^{-1}. Rewrite the CDS equation (6) as:

$$0 = 2X_i\tilde{\mathbb{E}}\left[\beta^{\tau_i - t}\mathbf{1}_{\tau_i < T}\middle|\,\mathcal{G}_t\right] - c_i(t,T)\sum_{k=1}^{n}\beta^{6s}\tilde{\mathbb{E}}\left[\mathbf{1}_{\tau_i \geq t+6s}\middle|\,\mathcal{G}_t\right].$$

[5]That is, I assume that the firm managers choose the debt level and the corresponding coupon payment to take optimal advantage of the debt tax shield as in Duffie and Lando (2001).

Recall that the expectation functions $\tilde{\mathbb{E}}\left[\beta^{\tau_i - t}\mathbf{1}_{\tau_i < T}\middle|\,\mathcal{G}_t\right]$ and $\tilde{\mathbb{E}}\left[\mathbf{1}_{\tau_i \geq t+6s}\middle|\,\mathcal{G}_t\right]$ are calculated using the worst-case likelihood and, thus, depend on investors' preferences toward ambiguity. Thus, to estimate θ_d^{-1} and θ_s^{-1}, I assume that the CDS rates are observed with a measurement error. In particular, assume that the observation equation is given by:

$$0 = 2X_i\tilde{\mathbb{E}}\left[\beta^{\tau_i - t}\mathbf{1}_{\tau_i < T}\middle|\,\mathcal{G}_t\right] - \hat{c}_i(t,T)\sum_{k=1}^{n}\beta^{6s}\tilde{\mathbb{E}}\left[\mathbf{1}_{\tau_i \geq t+6s}\middle|\,\mathcal{G}_t\right] + \eta_{iT,t},$$

where $\hat{c}_i(t,T)$ are the observed CDS rates and the vector of maturity-specific measurement errors $\eta_{it} = [\eta_{i1,t}, \ldots, \eta_{iT,t}]$ is normally distributed and i.i.d. across time and firms: $\eta_i \sim N(0, \Sigma_\eta)$. Taking the point estimates from the Gibbs sampling procedure as estimates of the reference model parameters, I make draws of the vector $\theta^{-1} = \left[\theta_d^{-1}\ \theta_s^{-1}\right]'$ using a Random Walk Metropolis algorithm with a flat prior. The accept/reject probability for the draws of θ^{-1} is the likelihood ratio of the CDS rates for all firms, at all available data points and for all available maturities.

The estimation procedure faces two identification concerns. The first is whether it is possible to separately identify the agents' degree of aversion to ambiguity about the information quality and to ambiguity about the underlying dynamics. Recall from above that, for longer maturities, the CDS spread is determined primarily by agents' beliefs about the future states while for shorter maturities the uncertainty about the current state also plays an important role. Recall further that the magnitude of the worst case distortion to the posterior distribution of the current state is determined by the degree of aversion to ambiguity about the information quality while the magnitude of the worst-case distortion to the transition probability matrix is determined by the degree of aversion to ambiguity about the underlying dynamics. Thus, considering CDS of different maturities allows us to separately identify the agents' preferences toward the two sources of ambiguity.

The second concern relates to differentiating between time variation in agents' attitudes toward ambiguity (as measured by θ_d^{-1} and θ_s^{-1}) and time variation in the amount of ambiguity in the economy (as measured by ϵ^1 and ϵ^2). Within a given sample, only the variation in one of these two quantities can be identified, with the other kept constant for the purposes of estimation. To circumvent this problem, I conduct three different estimations of the parameters θ_d^{-1} and θ_s^{-1} using different data sub-periods: the pre-crisis period (January 2004 – July 2007), the period from the start of the crisis to the bailout of Bear Stearns (August 2007 – March 2008) and the period from the bailout of Bear Stearns to the liquidation of Lehman Brothers (April 2008 – September 2008). While the estimation assumes that the agents' attitudes

toward ambiguity are constant within each subsample, examining the differences in the estimates on different sub-periods informs us about time variation in agents' attitudes between the sub-periods, separating it from time variation in the amount of ambiguity.

The results of these estimations are presented in Table 1. Notice that the estimation suggests that the investors in this market exhibit a higher degree of aversion to ambiguity about the filter distribution than to ambiguity about the underlying dynamics. Notice further that the three different periods do not yield significantly different estimates of θ_d^{-1} and θ_s^{-1}, suggesting that the investors attitudes toward model misspecification do not change significantly during the crisis period. Instead, the observed changes in CDS rates are due to the amount of ambiguity faced by agents in the markets and to shifts in what kind of misspecification the agents are more concerned about. Finally, consider the model error detection probabilities reported with the estimates. The model error detection probability measures the probability of making a mistake (that is, interpreting the data as being generated by the wrong model) in trying to distinguish between the reference and the misspecified model in a sample of the same size as the observed data (275 observations). Thus, for example, there's a 23% probability of making a mistake in distinguishing between the reference model and the worst-case model implied by the pre-crisis estimates of the ambiguity aversion parameters. Overall, the estimated detection error probabilities are similar to those found in the previous literature. For example, Barillas et al. (2009) find detection error probabilities of around 20% to be reasonable.

Consider now the model-implied time series evolution of credit spreads. Since the estimate of θ^{-1} does not change during the crisis, I use the pre-crisis estimates of θ_d^{-1} and θ_s^{-1} to compute the model-implied CDS spreads. Table 2 presents the observed 5 year CDS spreads for the financial institutions at five dates of interest – before the start of the crisis, July 2007, at the start of the crisis in August 2007, after the bailout of Bear Stearns in March 2008, after the liquidation of Lehman Brothers in September 2008 and after the introduction of TARP in October 2008 – together with the model-implied CDS rates at these dates. Notice first that the model-implied CDS rates follow the observed pattern of increasing during the financial crisis. Further, for most institutions, the implied CDS rates match the levels of CDS spreads over time, although the performance of the estimated model worsens after the liquidation of Lehman Brothers.

Recall that the estimates of the ambiguity aversion parameters θ_d^{-1} and θ_s^{-1} do not change significantly over the crisis. What then generates the time series behavior of CDS rates that we observed during the crisis? In this paper, I argue that the observed rapid increases in CDS spreads are driven by changes in the amount of

total ambiguity faced by investors in the market, as well as changes in how this total amount of ambiguity is decomposed into the amount of ambiguity about the underlying dynamics and the amount of ambiguity about the filter distribution. In Fig. 2, I plot the time series evolution of ambiguity, as measured by the conditional relative entropy between the reference and the worst-case likelihoods[6], during the crisis. Notice that at the beginning of the crisis (after the BNP Paribas announcement in August 2007), only ambiguity about the filter distribution increased. Intuitively, at the beginning of the crisis, although investors observed the negative reports released by financial institutions, since no defaults of major institutions occurred, they could interpret these new signals as informing about the quality of information and not about the quality of the assets on the balance sheets of these institutions. Contrast this with the changes to the amount of ambiguity after the bailout of Bear Stearns in March 2008 and after the liquidation of Lehman Brothers in September 2008. In these cases, it is the amount of ambiguity about the underlying dynamics that increases. Unlike the start of the crisis, investors now observed default by major financial institutions, which increased their uncertainty about the quality of assets held by financial institutions. Notice also that, although the amount of total entropy decreased a little after the introduction of TARP in October 2008, it was still higher than before the start of the crisis, which is consistent with the Caballero and Krishnamurthy (2008) intuition that ambiguity increased during the crisis.

To further understand the evolution of asset prices during the crisis, consider the time series evolution of expected time to default of the financial institutions. Table 3 presents the expected time to default for the financial institutions at five dates of interest – before the start of the crisis (July 2007), at the beginning of the crisis in August 2007, after the bailout of Bear Stearns in March 2008, after the liquidation of Lehman Brothers in September 2008 and after the introduction of TARP in October 2008 – together with the percentage change in the time to default relative to the previous month. The initial BNP Paribas announcement in August 2007 lead to a decrease in the expected time to default for only Bank of America, Goldman Sachs and Morgan Stanley. This coincides with the intuition that investors interpreted the announcements by financial institutions at the beginning of the crisis as increasing the amount of ambiguity about the filter distribution, and not necessarily indicative of the quality of assets held by the financial institutions. More interesting is the fact that the bailout of Bear Stearns in March 2008 lead to a decrease in the expected time to default for only some firms. In contrast, the liquidation of Lehman Brothers in September 2008 decreased the expected time

[6]See the definition of conditional relative entropy in eq. (10)

to default for all institutions except Goldman Sachs and Morgan Stanley. Thus, while the bailout of Bear Stearns induced contagion effects for some of the financial institutions, the effect was not as widespread as that induced by the liquidation of Lehman Brothers.

Compare this to the evolution of expected times to default under the reference model, presented in Panel B of Table 3. Notice first that, under the reference model, the expected time to default is longer at all dates than that under the misspecified model. Intuitively, the misspecification-averse agent perceives the probability of default next period to be greater than under the reference model, decreasing the expected time to default. Next, consider the time evolution of the expected times to default. The initial BNP Paribas announcement leads to a slight (0.03%) decrease in the expected time to default for each of the institutions. As under the misspecified model, the decrease in the expected time to default is much greater after the bailout of Bear Stearns. Notice, however, that the decrease under the reference model is greater than under the misspecified model. Intuitively, since the misspecification-averse agent already has more pessimistic views of the future, observing the bailout of Bear Stearns did not have as a large of an impact on her beliefs as it did on the beliefs under the reference model. Similarly, the decrease under the reference model after the liquidations of Lehman Brothers is larger than under the misspecified model. After introduction of TARP in October 2008, however, the increase in the expected time to default under the reference model is much smaller than under the misspecified model.

Finally, consider the implied time series evolution of equity prices during the crisis. Fig. 3 plots the observed evolution of equity prices together with the model-implied evolution. Although we cannot hope to match the level of equity prices since the model of firm earnings in this paper is extremely simplistic, the model should be able to match the observed movements in equity prices. Comparing the model-implied evolution to the true evolution of equity prices, we see that the model-implied equity prices lag the observed equity prices. This is not surprising since the signals used to construct the time series evolution of conditional probabilities are backward-looking; for example, the Case-Shiller 10 Index is constructed using observations over the previous three months. The model is able to capture the overall downward trend of equity prices during the crisis and especially well the sharp drop in equity prices after the bailout of Bear Stearns and after the liquidation of Lehman Brothers.

5. Conclusion

The rapid increases in CDS spreads observed during the recent financial crisis are puzzling from the standpoint of standard structural credit risk models, espe-

cially since initial increases in CDS spreads were not accompanied by decreases in equity prices. Introducing misspecification concerns to an incomplete information version of the Black and Cox (1976) model of credit spreads exacerbates the imperfect information problem faced by the representative agent. This leads to an increased level of default swap spreads overall and greater sensitivity of CDS spreads to bad news. The misspecification-averse agent perceives the probability of default next period to be higher than under the reference model, increasing CDS spreads and decreasing expected time to default. Observing a bad signal not only increases the conditional probability of being in a low-payoff state in the current period but also increased the perceived probability of default in the next period.

Estimating this model using data on nine major financial institutions, I find that, while agents' preferences toward model misspecification did not change, the total amount of ambiguity in the economy increased during the crisis. Furthermore, the relative contributions of ambiguity about the underlying dynamics and of ambiguity about the signal quality shifted during the crisis, causing a disconnect between credit spreads and equity prices in the early part of the crisis. In particular, the initial BNP Paribas announcement in August 2007 led to an increase in ambiguity about the signal quality, increasing CDS spreads but not decreasing equity prices. The bailout of Bear Stearns in March 2008 and the liquidation of Lehman Brothers in September 2008 on the other hand led to an increase in the amount of ambiguity about the underlying dynamics, increasing CDS spreads and decreasing equity prices.

While the model considered in the paper is geared toward explaining observed increases in CDS spreads, similar intuition could be used to explain observed changes to prices of collateralized debt obligations (CDOs) and other complex securities. In fact, since arguably CDOs have a more complicated underlying structure than default swaps, model misspecification concerns would be even more relevant in pricing these securities. A formal treatment of this problem, however, is left for future research.

Anderson, E. W., Hansen, L. P., Sargent, T. J., 2003. A quartet of semigroups for model specification, robustness, prices of risk and model detection. Journal of the European Economic Association 1 (1), 68–123.

Anderson, E. W., Hansen, L. P., Sargent, T. J., 2012. Small noise methods for risk-sensitive/robust economies. Journal of Economic Dynamics and Control 36 (4), 468–500.

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

Barillas, F., Hansen, L., Sargent, T., 2009. Doubts or variability? Journal of Economic Theory 144 (6), 2388–2418.

Bhamra, H., Kuehn, L.-A., Strebulaev, I., 2010. The Levered Equity Risk Premium and Credit Spreads: A Unified Framework. Review of Financial Studies 23 (2), 645–703.

Black, F., Cox, J., 1976. Valuing corporate securities: Some effects of bond indenture provisions. Journal of Finance 31 (2), 351–367.

Caballero, R. J., Krishnamurthy, A., 2008. Collective risk management in a flight to quality episode. Journal of Finance 63, 2195–2230.

Campbell, J., Cochrane, J., 1999. By force of habit: A consumption-based explanation of aggregate stock-market behavior. Journal of Political Economy 107, 205–251.

Campi, M., James, M., 1996. Nonlinear discrete-time risk-sensitive optimal control. International Journal of Robust and Nonlinear Control 6, 1–19.

Chen, H., 2010. Macroeconomic conditions and the puzzles of credit spreads and capital structure. Journal of Finance 65, 2171–2212.

Chen, L., Collin-Dufresne, P., Goldstein, R., 2009. On the relation between the credit spread puzzle and the equity premium puzzle. Review of Financial Studies 22, 3367–3409.

Chen, Z., Epstein, L. G., 2002. Ambiguity, risk and asset returns in continuous time. Econometrica 70, 1403–1443.

Drechsler, I., 2010. Uncertainty, time-varying fear and asset prices, Working Paper.

Duffie, D., Lando, D., 2001. Term structures of credit spreads with incomplete accounting information. Econometrica 69, 633–664.

Epstein, L. G., Miao, J., 2003. A two-person dynamic equilibrium under ambiguity. Journal of Economic Dynamics and Control 27, 1253–1288.

Epstein, L. G., Schneider, M., 2008. Ambiguity, information quality and asset pricing. Journal of Finance LXIII (1), 197–228.

Epstein, L. G., Zin, S., 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57, 937–969.

Gagliardini, P., Porchia, P., Trojani, F., 2009. Ambiguity aversion and the term structure of interest rates. Review of Financial Studies 22, 4157–4188.

Hansen, L., Sargent, T., 1995. Discounted linear exponential quadratic gaussian control. IEEE Trans. Autom. Control 40, 968–971.

Hansen, L., Sargent, T., 2005. Robust estimation and control under commitment. Journal of Economic Theory 124 (2), 258–301.

Hansen, L., Sargent, T., 2007. Recursive robust estimation and control without commitment. Journal of Economic Theory 136, 1–27.

Hansen, L., Sargent, T., Tallarini Jr., T., 1999. Robust permanent income and pricing. Review of Economic Studies 66, 873–907.

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

Leland, H., 1994. Corporate debt value, bond covenants and optimal capital structure. Journal of Finance 49 (4), 1213–1252.

Liu, J., Pan, J., Wang, T., 2005. An equilibrium model of rare-event premia and its implication for option smirks. Review of Financial Studies 18, 131–164.

Trojani, F., Vanini, P., 2002. A note on robustness in merton's model of intertemporal consumption and portfolio choice. Journal of Economic Dynamics and Control 26 (3), 423–435.

Trojani, F., Vanini, P., 2004. Robustness and ambiguity aversion in general equilibria. The Review of Finance 2, 423–435.

Uppal, R., Wang, T., 2003. Model misspecification and under-diversification. Journal of Finance LIX, 2465–2486.

	Pre-crisis (01/04–07/07)	Pre- Bear Stearns (08/07–03/08)	Pre- Lehman Brothers (04/08–09/08)
θ_d^{-1}	0.23	0.21	0.29
	(0.017)	(0.028)	(0.018)
θ_s^{-1}	0.64	0.65	0.69
	(0.021)	(0.029)	(0.019)
MSE	2.63	4.04	8.51
Detect. Prob.	0.23	0.21	0.23

Table 1: Estimates of investors' degree of aversion to ambiguity about the underlying dynamics, θ_d^{-1}, and degree of aversion to the filter distribution, θ_s^{-1}. Half-width of the 95% confidence intervals are reported in parentheses. Mean Squared Error (MSE) is the average squared error between the observed CDS rates and the model-implied CDS rates for all institutions, maturities, and observations. The model detection error probability (Detect. Prob.) measures the probability of not distinguishing between the reference model and the implied worst-case model in a sample of size 275 months. The reference model parameters are estimated using monthly observations from May 1987 through July 2007, inclusively. Draws are made using the Metropolis-Hastings procedure of Section 4.2.

	BAC	BSC	Citi	GS	JPM	MER	MS	WB	WFC
Jul 31 07	36.2	161.7	37.2	81.2	55	74.7	75.2	39.5	35.9
	29.6	51.3	34.9	71.2	54.7	73.8	73.2	39.3	38.7
Aug 31 07	39.7	135.7	45.5	68.8	45.4	71.7	68.8	39.4	35
	22.1	128.7	41.4	36.5	43.2	72.2	64.7	32.3	32.2
Mar 31 08	86.8	122.7	138.2	115	87.5	195.8	153.9	142.8	80.8
	65.1	157.0	128.62	133.9	85.3	109.7	93.4	83.2	90.9
Sep 30 08	170	143.3	301.7	452.5	143.8	410.8	1022	385.8	170
	102.7	144.5	486.4	482.8	152.6	380.03	317.3	386.5	138.1
Oct 31 08	133.1	120.2	197.6	313.3	119.9	216.2	413.3	121	97.2
	133.3	161.3	178.8	380.7	146.6	203.5	416.26	131.3	93.1

Table 2: Observed 5 year CDS rates and model-implied 5 year CDS rates at five different dates. The reference model parameters are estimated using the Gibbs sampling procedure of Online Appendix for the time period May 1987 through July 2007; the degree of aversion to ambiguity about the underlying dynamics, θ_d^{-1}, and the degree of aversion to ambiguity about the filter distribution, θ_s^{-1}, are estimated using the Metropolis-Hastings algorithm of Section 4.2 for the time period January 2004 through July 2007.

Panel A: Misspecified Model

	BAC	BSC	Citi	GS	JPM	MER	MS	WB	WFC
Jul 31 07	5.15	9.54	10.04	13.32	11.65	10.80	13.34	11.02	10.96
	100.43	-18.79	-25.16	256.87	-9.95	-32.18	47.10	-26.28	-28.05
Aug 31 07	2.61	11.74	13.42	3.77	13.01	16.03	9.11	15.03	15.32
	-49.44	23.04	33.73	-71.71	11.64	48.39	-31.73	36.44	39.81
Mar 31 08	4.81	0.00	16.22	5.12	11.64	13.42	9.08	12.83	13.01
	5.11	-100.00	19.85	37.19	-9.44	-15.05	0.53	-13.45	-13.90
Sep 30 08	4.27	–	9.55	10.31	11.50	10.47	13.77	10.74	10.66
	20.39		-42.75	155.77	-0.53	-23.03	59.41	-16.98	-18.83
Oct 31 08	5.00	–	13.86	4.55	10.44	11.81	8.25	11.34	11.49
	17.20		45.09	-55.83	-9.24	12.77	-40.09	5.61	7.72

Panel B: Reference Model

	BAC	BSC	Citi	GS	JPM	MER	MS	WB	WFC
Jul 31 07	16.47	16.47	16.47	16.47	16.47	16.47	16.47	16.47	16.47
	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
Aug 31 07	16.46	16.46	16.46	16.46	16.46	16.46	16.46	16.46	16.46
	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03
Mar 31 08	11.23	0.00	11.26	11.23	11.21	11.19	11.21	11.22	11.21
	-31.74	-100.00	-31.63	-31.79	-31.88	-31.98	-31.90	-31.81	-31.87
Sep 30 08	11.23	–	11.25	11.25	11.26	11.22	11.21	11.22	11.22
	-28.05		-27.97	-27.99	-27.92	-28.12	-28.18	-28.13	-28.15
Oct 31 08	11.32	–	11.35	11.34	11.35	11.32	11.31	11.32	11.31
	0.81		0.85	0.84	0.79	0.84	0.85	0.87	0.85

Table 3: Expected time to default and the percentage change in the expected time to default relative to previous month for different financial institutions. Panel A: expected time to default perceived by the misspecification-averse agent. Panel B: expected time to default under the reference model. The reference model parameters are estimated using the Gibbs sampling procedure of Online Appendix for the time period May 1987 through July 2007; the degree of aversion to ambiguity about the underlying dynamics, θ_d^{-1}, and the degree of aversion to ambiguity about the filter distribution, θ_s^{-1}, are estimated using the Metropolis-Hastings algorithm of Section 4.2 for the time period January 2004 through July 2007.

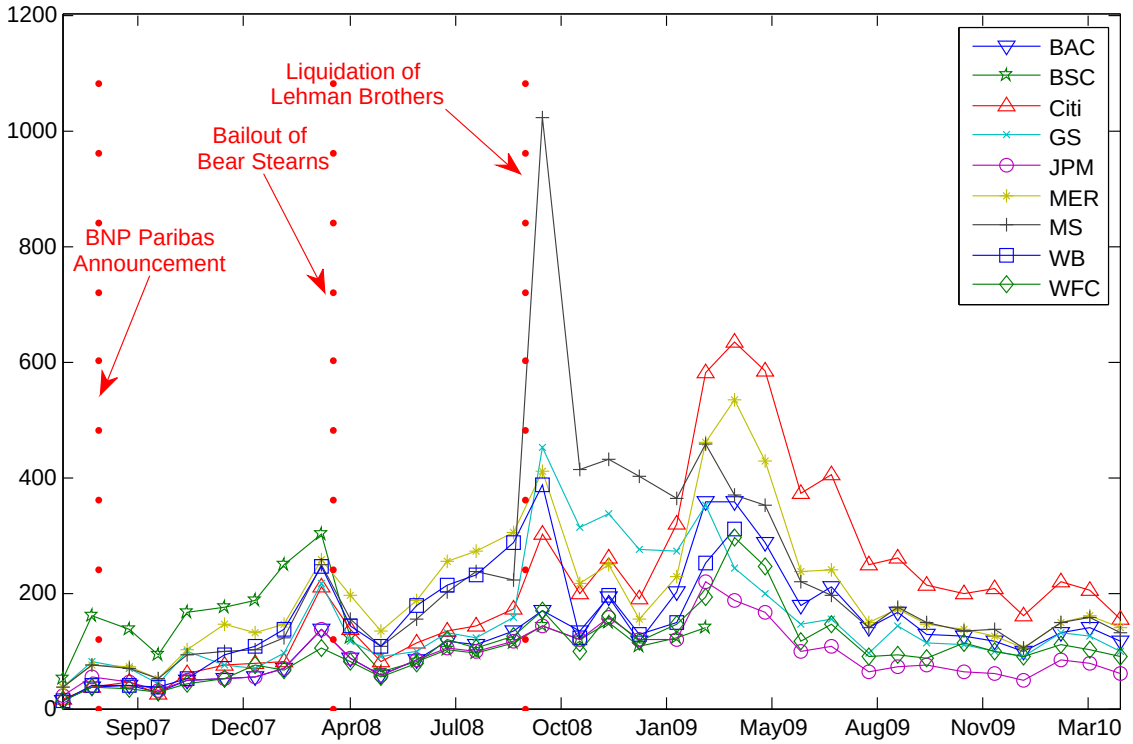

Figure 1: Evolution of the five year CDS spreads for financial institutions over the course of the crisis. Data source: Datastream

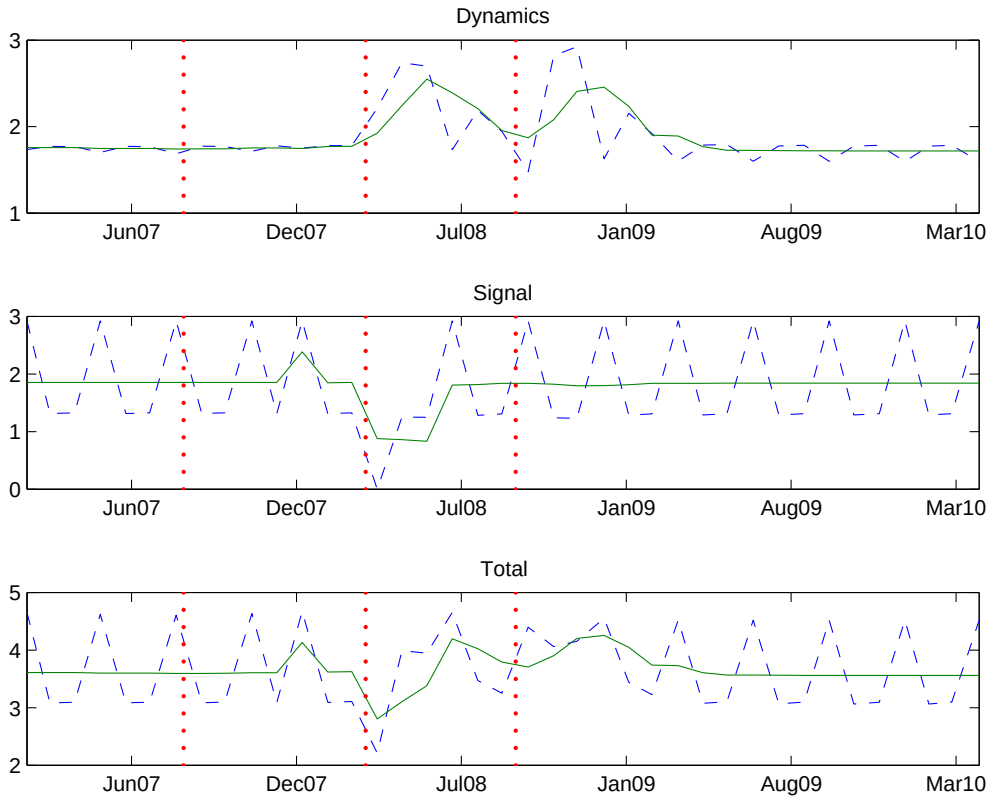

Figure 2: Components of relative entropy between the reference and misspecified models over time. Upper panel: entropy due to misspecification of the joint signals and states dynamics; central panel: entropy due to misspecification of the current period conditional probability; lower panel: total entropy. The solid line in each panel is the three month moving average of the corresponding entropy measure. The three vertical lines correspond, respectively, to the BNP Paribas announcement in August 2007, the bailout of Bear Streans in March 2008 and the liquidation of Lehman Brothers in September 2008. The reference model parameters are estimated using the Gibbs sampling procedure of Online Appendix for the time period May 1987 through July 2007; the degree of aversion to ambiguity about the underlying dynamics, θ_d^{-1}, and the degree of aversion to ambiguity about the filter distribution, θ_s^{-1}, are estimated using the Metropolis-Hastings algorithm of Section 4.2 for the time period January 2004 through July 2007.

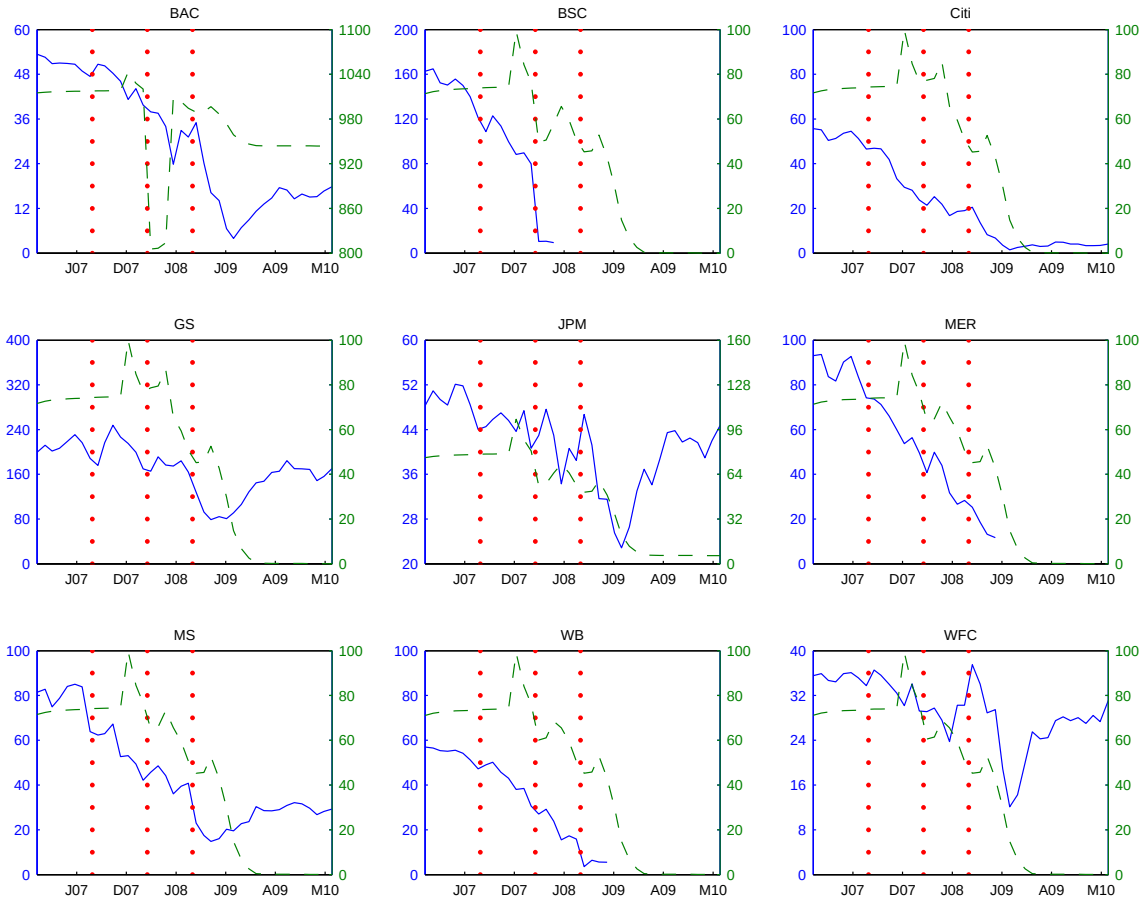

Figure 3: Value of equity under the misspecified model (dashed line, right-hand scale) and the observed value of equity (solid line, left-hand scale) of the different financial institutions over time. The three vertical lines correspond, respectively, to the BNP Paribas announcement in August 2007, the bailout of Bear Streans in March 2008 and the liquidation of Lehman Brothers in September 2008. The reference model parameters are estimated using the Gibbs sampling procedure of Online Appendix for the time period May 1987 through July 2007; the degree of aversion to ambiguity about the underlying dynamics, θ_d^{-1}, and the degree of aversion to ambiguity about the filter distribution, θ_s^{-1}, are estimated using the Metropolis-Hastings algorithm of Section 4.2 for the time period January 2004 through July 2007.