Real Credit Cycles*

VERY PRELIMINARY AND INCOMPLETE

Pedro Bordalo[†] *University of Oxford*

Nicola Gennaioli[‡] *Universita Bocconi*

Andrei Shleifer[§] *Harvard University*

Stephen J. Terry[¶] *Boston University*

July 2020

Abstract

Recent empirical work has revived the Minsky hypothesis of boom-bust credit cycles driven by fluctuations in investor optimism. To quantitatively assess this hypothesis, we incorporate diagnostic expectations into an otherwise standard business cycle model with heterogeneous firms and risky debt. Diagnostic expectations are a psychologically founded, forward-looking model of belief formation that captures over-reaction to news. We calibrate the diagnosticity parameter using microdata on the forecast errors of managers of listed firms in the US. The model generates countercyclical credit spreads and default rates, while the rational expectations version generates the opposite pattern. Diagnostic expectations also offer a good fit of three patterns that have been empirically documented: systematic reversals of credit spreads, systematic reversals of aggregate investment, and predictability of future bond returns. Crucially, diagnostic expectations also generate a strong fragility or sensitivity to small bad news after steady expansions. The rational expectations version of the model can account for the first pattern but not the others. Diagnostic expectations offer a parsimonious account of major credit cycles facts, underscoring the promise of realistic expectation formation for applied business cycle modeling.

*We thank audiences at FRB Minneapolis, Yale Finance, Harvard Business School Finance, NBER Behavioral Macro, Stanford SITE Uncertainty, NHH, the European Central Bank, Boston University, and Harvard Economics for valuable comments and suggestions.

[†]pedro.bordalo@sbs.ox.ac.uk

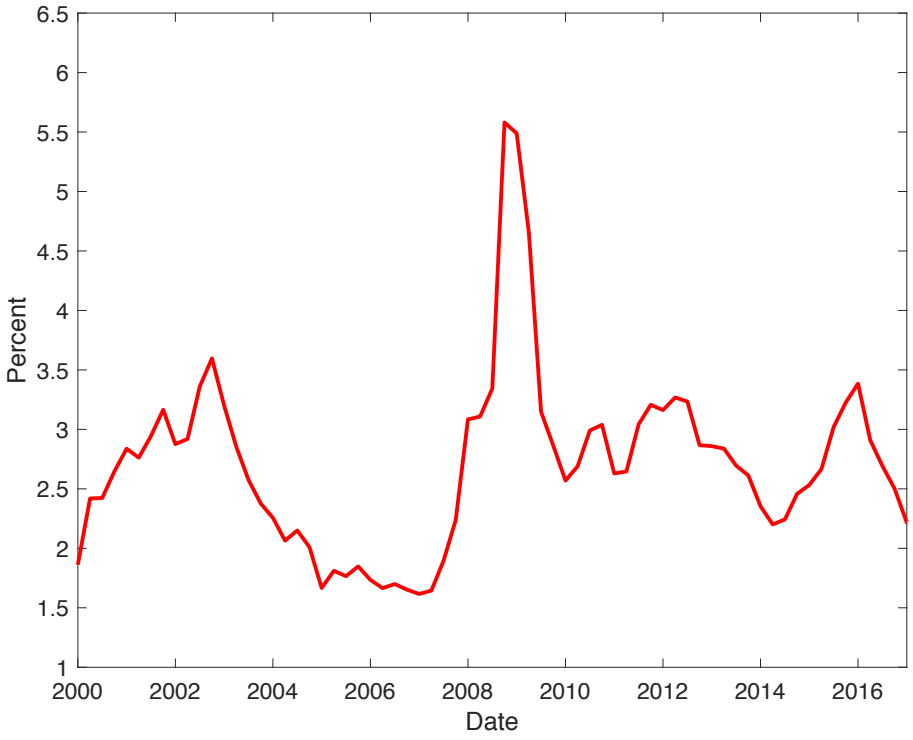
[‡]nicola.gennaioli@unibocconi.it

[§]shleifer@fas.harvard.edu

[¶]stephent@bu.edu

Before the summer of 2007, the spread between risky corporate bonds and safe interest rates was unusually low. Within one year, after the tremors of August 2007 and especially after the Lehman collapse on September 2008, the spread reversed to historical highs (see Figure 1). The earliers booms in corporate credit and investment also reversed as the economy moved into the Great Recession. These boom bust dynamics in credit spreads, leverage, and economic activity are generally observed around financial crises (Borio and Lowe, 2002b,a; Schularick and Taylor, 2012; Jordà et al., 2013; Krishnamurthy and Muir, 2016), but also, in a less dramatic form, during normal times in the US and other economies (Mian et al., 2017; López-Salido et al., 2017).

Figure 1: Interest Rate Spreads Fell before the Financial Crisis



Notes: The figure plots the average spread of the Moody's BAA corporate bond yield relative to the 10-Year Treasury yield, both in annualized terms, from 2000 through 2017 at quarterly frequency.

Two patterns in the data suggest that non-rational expectations may play a role in these events. First, there is systematic return predictability: After credit booms, the realized excess returns of corporate bonds and bank stocks are systematically disappointing and even negative (Baron and Xiong, 2017; Greenwood and Hanson, 2013). Second, there is direct evidence from expectations data. When spreads are low, credit analysts' forecasts of future spreads are too low as well (Bordalo et al., 2018), and stock analysts' forecasts about the future profitability of risky firms are too optimistic (Gulen et al., 2019). These facts suggest that during good times market

participants may neglect the buildup of risk, causing credit spreads to be low, and credit volumes and investment to be high. These belief errors are subsequently corrected, causing systematic reversals of spreads, investment, and economic activity, in line with the early analysis of Minsky (1977).

This evidence raises two questions. First, can realistic departures from rational expectations improve the ability of standard business cycle models to account for these facts? Second, how far can non-rational expectations go in accounting quantitatively for instability in credit markets and the macroeconomy? To address these questions, we incorporate diagnostic expectations into an otherwise standard model of total factor productivity (TFP) driven business cycles.

We study a heterogeneous firms model (Bachmann et al., 2013; Khan and Thomas, 2008) in which the productivity of each firm is subject to an idiosyncratic and a common component, both of which follow AR(1) processes. Each firm optimally decides whether or not to default on debt, how much labor to hire, how much to invest with adjustment costs, and how many one-period bonds to issue. Although we eventually extend the model to general equilibrium in Section 8, our main formulation is partial equilibrium for tractability: funds are provided by deep pocketed risk neutral lenders, and labor is infinitely elastic at an exogenous fixed wage. The key non-standard ingredient is that both firms and lenders hold diagnostic rather than rational expectations.

Diagnostic expectations (Bordalo et al., 2018) are a model of belief formation built on the representativeness heuristic of probabilistic judgments introduced by Kahneman and Tversky (1972), and ultimately based on the psychology of human memory (Kahana, 2012; Bordalo et al., Forthcominga; Gennaioli and Shleifer, 2010). They have been shown capable of accounting for survey evidence on the expectations of financial analysts (Bordalo et al., 2019b) and macroeconomic forecasters (Bordalo et al., Forthcomingb), and have been used to shed light on financial fragility (Bordalo et al., 2018; Gennaioli and Shleifer, 2018). Diagnostic expectations are formed through a selective recall process that captures a form of over-reaction to current news, so that individuals become too optimistic after good news and too pessimistic after bad news. In forecasting future macro TFP A_{t+1}, this process yields the following representation of beliefs that we derive in Section 2:

$$\mathbb{E}_t^{\theta}\left(A_{t+1}\right) = \mathbb{E}_t\left(A_{t+1}\right) + \theta\left[\mathbb{E}_t\left(A_{t+1}\right) - \mathbb{E}_{t-1}\left(A_{t+1}\right)\right] \tag{1}$$

where $\mathbb{E}_t\left(\cdot\right)$ is the rational expectation at time t and $\theta \geq 0$ is the diagnosticity parameter measuring the extent of over-reaction to news. The model nests rational expectations for $\theta = 0$. Equation (1) illustrates the so called "kernel of truth property," whereby diagnostic expectations exaggerate true patterns in the data. This feature yields two important properties. First, deviations from rationality are disciplined by the data generating process, which pins down $\mathbb{E}_t\left(\cdot\right)$, and are summarized by the single parameter θ. Second, deviations from rationality are forward

looking and react to regime changes, so they are robust to the Lucas (1976) critique of adaptive expectations.

We calibrate the model by relying on conventional estimates for some parameters and matching moments from microdata on firm level profitability, leverage and investment. To quantify the degree of diagnosticity θ we use the model to match the predictability of managerial errors in forecasting their firm profits, which are computed using microdata obtained from the Compustat and Institutional Brokers' Estimate System (IBES) databases. As shown in Gennaioli et al. (2016), managers are too optimistic about future profits when current profits are high and too pessimistic when current profits are low. To match this fact, the calibration sets $\theta \sim 1$, which is in the ballpark of estimates obtained using different data and in different domains (Bordalo et al., 2019b, Forthcomingb). The calibrated diagnostic expectations (DE) model significantly improves the fit relative to the associated rational expectations (RE) model with $\theta = 0$. Sizable departures from rational expectations are needed to account for firm level behavior.

When we simulate the model, we show that diagnosticity has significant macroeconomic effects. Consistent with the evidence, the DE model features countercyclical credit spreads. This is due to the large shifts in credit supply by diagnostic investors, who underestimate default risk in good times, leading to low spreads. In fact, in the RE model spreads are procyclical, mostly moved by demand for capital. Because diagnostic expectations on average revert back to rationality, our DE model generates the following four patterns:

1. More macroeconomic volatility than RE for the same fundamental shocks.

2. Financial crises – defined as periods with large increases in credit spreads – are preceded by low spreads as in Figure 1, and are triggered by the slowdown of TFP growth, not by bad shocks.

3. Systematic reversals of spreads and investment: a reduction in the current spread predicts a hike in the spread and a reduction in investment in the future, as documented by López-Salido et al. (2017).

4. Predictability of future bond returns: a reduction in the current spread predicts disappointing realized bond returns, as documented by Greenwood and Hanson (2013).

5. A large decline in credit growth during financial crises and the recent Great Recession episode.

The RE model can partly account for reversals in spread growth (both during large crises and during normal fluctuations) due to fundamental mean reversion in productivity, but it cannot account for the other facts. In particular, the RE model does not generate large declines in credit

spreads before the crisis (Krishnamurthy and Muir, 2016). In the DE model this phenomenon is natural due to: i) large pre-crisis expansion of the supply of capital by over-optimistic lenders, and ii) high borrowing and leverage at firms that enhances fragility when optimism wanes. The same mechanism explains why in the DE model crises occur at the end of booms, and not after bad times.

The RE model also cannot account for the association between low spreads today, and low investment and low realized bond returns in the future. In the RE model, realized returns are on average equal to lenders' required return, which is constant. In the DE model, return predictably is due to under-pricing of credit risk in good times and over-pricing of credit risk in bad times. Likewise, the DE model predicts systematic reversal in investment due to the systematic disappointment of expectations after good times and the systematic improvement of expectations after bad times.

To assess the relative role of expectations errors by lenders and borrowers, we also simulate a model in which only firms are diagnostic. In this model, crisis dynamics are qualitatively similar but more muted. Spreads increase during crises by only around half of their amount in the baseline DE model, suggesting that credit supply shifts are an important part of the mechanism.

We conclude the analysis by studying the extent to which our model can quantitatively account for the 2008 US crisis. We fit an aggregate TFP sequence that allows our model to match the actual dynamics of US investment growth during 2007-2012. We then consider the implication of this TFP sequence for leverage and spreads both in the DE and RE models. We find that diagnosticity generates the 2008 crisis and Great Recession with a fairly mild negative TFP shock of -1.5%. The RE model can neither match the collapse in investment in 2008 nor its subsequent fast recovery, which the DE model sees as a reversal of excess pessimism in 2008. Furthermore, only the DE model accounts for the dynamics of spreads and leverage.

In sum, our analysis conveys three messages. First, diagnostic expectations naturally generate boom-bust credit cycles that transmit to the real economy. Second, these dynamics are quantitatively sizable, helping account for real world phenomena such as predictable changes in credit and real markets. More generally, we show that psychologically founded models of non-rational expectation formation can be used in conventional business cycle models, can be disciplined with microdata, and can improve the ability of workhorse macroeconomic models to account for important phenomena such as credit cycles. They can be used to quantitatively assess macroeconomic outcomes in the same manner as the rich mechanisms studied in the heterogeneous firms literatures on adjustment costs (Khan and Thomas, 2013; Bachmann et al., 2013), uncertainty (Bloom, 2009; Christiano et al., 2014), firm-level financial frictions (Gilchrist and Zakrajšek, 2012; Gilchrist et al., 2014; Alfaro et al., 2018; Khan and Thomas, 2013), financial intermediation (He and Tian, 2013), and financial dynamics more generally (Brunnermeier and Sannikov, 2014). Barrero (2018)

analyzes behavioral beliefs in a heterogeneous firms model as well, although our study on business cycles and credit dynamics differs from his focus on steady-state misallocation.

Our paper is related to several literatures. First, a classic literature studies financial frictions, in the form of collateral constraints, as mechanisms that amplify economic shocks (Bernanke and Gertler, 1989; Bernanke et al., 1999; Kiyotaki and Moore, 1997). This work sometimes features "financial shocks" to collateral constraints as a way to capture non-fundamental disruption in financial markets. Gu et al. (2013) study a model of credit in which collateral constraints are endogenized via the threat of exclusion from financial markets in case of default. They show that in this setting, depending on investors' beliefs, there may be equilibria featuring endogenous credit cycles. Relative to these papers, we emphasize the importance of departures from rational expectations, particularly in creating shifts in the supply of capital, and endogenize "financial shocks" through the predictable correction or reversal of expectations errors.

Another growing literature studies financial crises and credit cycles. Arellano et al. (2019) analyze the 2008 from the vantage point of financial frictions, and introduce uncertainty shocks to account for a range of features of the crisis, including the decline in debt purchases, output and labor during the Great Recession, despite the relative stability of total factor productivity. In their model, uncertainty shocks remain a primitive. Here we focus on non rational expectations, which helps to explain predictability of forecast errors and of credit market conditions, and to generate substantial credit cycles and leverage dynamics even in the absence of uncertainty shocks.

Third, several studies of financial fragility consider intermediary leverage and bank runs (Brunnermeier et al., 2012; He and Krishnamurthy, 2019; Eggertsson and Krugman, 2012; Guerrieri and Lorenzoni, 2017). Here we abstract from intermediary leverage, but as we argue in Section 6 this factor clearly plays a key role in crises. Relative to this work, our approach helps explain why crises often come as sudden reversals after booms and in this sense are predictable (Baron and Xiong, 2017; Baron et al., 2019; López-Salido et al., 2017). Maxted (2019) introduces diagnostic expectations into He and Krishnamurthy (2019)'s model, and shows that such expectations can further exacerbate the instability arising from intermediary leverage.

Another approach to crises and macro financial fragility or cycles is more behavioral and emphasizes the importance of beliefs. Bordalo et al. (2018) offer a stylized model of credit cycles with diagnostic expectations. Greenwood et al. (2019) build a model in which credit markets extrapolate from recent default history, so that crises are slow moving. Fostel and Geanakoplos (2014) and Simsek (2013) emphasize belief heterogeneity. Empirically, Bianchi et al. (2020) measure macro belief distortions using professional forecaster data. Relative to these papers, our main contribution is to introduce diagnostic expectations into a workhorse, quantitative macro model with a focus on heterogeneity and a quantified link to micro evidence. Bhandari et al. (2019) and Jaimovich and Rebelo (2007) also study the impact of belief distortions on business cycles

in quantitative macro models using distinct belief structures. Finally, a different approach to financial instability views reversals after booms as the result of slow reallocation of excess-capital toward more productive sectors Rognlie et al. (2018), fire sales (Shleifer and Vishny, 1992; Lorenzoni, 2008; Stein, 2012; Dávila and Korinek, 2017), or demand externalities (Farhi and Werning, 2016; Korinek and Simsek, 2016). This work adopts rational expectations, so it cannot account for return predictability and predictable expectations errors.

Section 2 introduces our notion of diagnostic beliefs and sketches a two-period model for intuition. Section 3 introduces our quantitative real business cycle model with diagnosticity. Section 4 describes our data and calibration approach. Section 5 evaluates the consequences of diagnosticity relative to the rational expectations model. Section 6 performs a model-based decomposition of the 2008 US financial crisis. Section 7 shows that the DE model entails state-dependent or nonlinear responses of investment. Section 9 concludes. Appendix A details our computational approach. Appendix B discusses the details of the microdata.

1 Firm-Level Credit Spreads and Forecast Errors

Credit spreads in Figure 1 alternate between periods of buoyant credit market conditions and sharp reversals, i.e., a credit cycle pattern appears prevalent at the macro level. Over the same period, beliefs or consensus forecasts of future spreads also fluctuate substantially. In good times with low spreads just before the financial crisis, consensus forecasts of credit spreads were overly optimistic and low, while forecasters expected overly pessimistic and high future spreads during the crisis itself (Bordalo et al., 2018). Taken together, the two patterns suggest that belief shifts might link to credit cycles. However, pushing beyond this suggestive time series evidence, in this section we make two empirical contributions based on disaggregated data.

First, we explore rich microdata on firm-level forecasts from the IBES manager guidance database. This panel dataset of forecasts records, for an individual firm-fiscal year, the prediction or guidance offered by the firm's management for their own company's profits or earnings over the next year. We exploit bundled forecasts, i.e., predictions made concurrently with the release of the current year's financials, in a sample also spanning the 1999-2017 period. The data from US public firms links to the Compustat database providing the standard firm financial information. Appendix B provides more information on the data and our sample. We then run a set of simple predictive regressions in Table 1 linking future forecast errors – realized minus predicted profits – to current firm-level financial outcomes. In the simplest version of full information rational expectations models, future forecast errors are fully unpredictable based on any information previously available to the firm. In contrast, if beliefs about firm-level fundamentals overreact to current conditions, displaying overoptimism during good times and undue pessimism during bad times,

Table 1: Predictable Forecast Errors

	(1)	(2)	(3)	(4)
	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$
Profits$_t$	-0.043**			
	(0.021)			
Investment$_t$		-0.455***		
		(0.065)		
Debt Issuance$_t$			-0.193**	
			(0.085)	
Forecast$_t$				-0.242***
				(0.030)
Firm Effect	X	X	X	X
Time Effects	X	X	X	X
Years	1999-17	1999-17	1999-17	1999-17
Firm-Years	9666	9666	9666	9666

Notes: The table reports estimates of specifications on the merged Compustat - IBES Guidance sample at the firm-fiscal year level. Profits are earnings, investment is tangible capital expenditures, debt issuance is end-of-period debt, and forecast errors are actual earnings minus manager forecast earnings at a 1-period horizon. All series are relative to firm tangible capital stocks. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, and ***=1% level. The standard deviation of future forecast errors is 0.822, the standard deviation of profits is 1.081, the standard deviation of investment is 0.241, the standard deviation of debt issuance is 0.283, and the standard deviation of forecasts is 0.939. For all series, 0.01=1%.

then future forecast errors should be negatively linked to current firm fundamentals. Each specification in Table 1 includes firm and time effects, isolating the variation we exploit to within-firm movements not shifting in a common manner across firms. Column 1 reveals that firms with higher profits today are systematically disappointed in future. In column 2 we show that firms investing more today issue overly optimistic forecasts. Column 3 reports that firms issuing more debt today are also disappointed ex-post. And in column 4 we see that firms making high forecasts today experience more negative surprises on average. Meaningful magnitudes are at work here. A firm with with an investment rate one standard deviation or 24% higher experiences on average a 24% $\times$ 0.455 $\approx$ 11% stronger disappointment in future.[1] To summarize, good times at firms are also systematically times with overly optimistic forecasts and lower future forecast errors.

Firm profit forecasts revert in a manner consistent with overreaction or extrapolation in Table 1. But does apparent forecast overreaction, and subsequent disappointment of overly optimistic firms, link to credit cycles at the firm level? In our second contribution, reported in Table 2, we

[1] A series of checks in Appendix B reveals robustness in the conclusions from Table 1 to a range of different sample choices.

Table 2: Linking Forecast Errors and Credit Spreads

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	$Return_t$	$Return_t$	$Return_t$	Δ $Spread_t$	Δ $Spread_t$	Δ $Spread_t$
Estimation	OLS	IV	Red. Form	OLS	IV	Red. Form
$Fcst.\ Error_t$	0.003**	0.104***		-0.005***	-0.088***	
	(0.001)	(0.037)		(0.001)	(0.029)	
$Spread_{t-2}$			0.248***			-0.212***
			(0.031)			(0.031)
First Stage		$Fcst.\ Error_t$			$Fcst.\ Error_t$	
$Spread_{t-2}$		2.394***			2.394***	
		(0.711)			(0.711)	
Years	2004-18	2004-18	2004-18	2004-18	2004-18	2004-18
Firm-Years	2410	2410	2410	2410	2410	2410
Firm Effects	X	X	X	X	X	X
Time Effects	X	X	X	X	X	X
First Stage F		11.33			11.33	

Notes: The table reports estimates of specifications on the merged Compustat - FISD/TRACE sample at the firm-fiscal year level. The top panel plots OLS, second-stage IV, and IV reduced-form estimates. The bottom panel, where relevant, reports first-stage estimates. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, and ***=1% level. The standard deviation of the bond return is 0.014, the standard deviation of spread growth is 0.024, the standard deviation of the forecast error is 0.433, and the standard deviation of the twice lagged spread level is 0.024. All series are in fractional units, so 0.01=1%.

show that the answer to this question is yes. To study firm credit spreads, we use a comprehensive dataset spanning essentially the universe of US corporate bonds from the early 2000s onwards. The Mergent Fixed Income Securities Database (FISD) contains issuance information on individual securities, while FINRA's Trade Reporting and Compliance Engine (TRACE) dataset contains detailed information on secondary market transactions from bond dealers. Combining the issuance and secondary market information allows for the calculation of firm-level bond spreads and returns. Linking to our dataset of IBES manager forecasts and Compustat information on firm financials, we can directly study the firm-level dynamics of firm beliefs, credit spreads, and bond returns for a panel spanning 2004-18 and over 400 large US public firms. Appendix B provides more information on the data and our sample.

We begin with two simple observations in Table 2. Columns 1 and 4 report OLS regressions of bond returns and credit spread growth on the forecast errors of firms in the same period. Firms with lower forecast errors – lower profits relative to expectations – experience lower bond returns and higher spread growth. However, this result would be entirely expected if current news drives both the credit pricing today and the current level of forecast errors. So we push further using an

IV specification to examine more restrictive variation.[2] We demonstrate in a strong common first stage in columns 2 and 5 that past credit spreads predict current forecast errors. In particular, firms with lower credit spreads in the past are more likely to see lower current forecast errors. The second stages in columns 2 and 5 reveal that exactly these firms also suffer credit market reversals – lower bond returns and sharp growth in their credit spreads – when their forecast errors decline. The predictions here prove to be quantitatively sizable, with a decline in a firm's forecast error of one standard deviation or 43% linked to an average coincident drop in bond returns of 43% $\times$ 0.104 $\approx$ 4.5% and spread growth of around 43% $\times$ 0.088 $\approx$ 3.8%. Columns 3 and 6 provide the reduced form of these IV regressions, directly linking lower past credit spreads to predictions of lower current bond returns and higher credit spread growth. We present all of these results in conservative specifications with both firm and time effects, so we emphasize that none of the patterns reflect simple heterogeneity or a particular macro episode with common shifts across firms.

To summarize, firms enjoying low credit spreads are likely to issue forecasts which are overly optimistic, followed by disappointing profits coincident with sharp credit market reversals. Motivated by this empirical evidence, in the next section we build a model laying out an environment in which microfounded overreaction or extrapolation in beliefs drives future reversion of forecast errors as well as credit market dynamics.

2 Diagnostic Expectations, Investment and Neglected Risk

We now introduce our model of Diagnostic Expectations and illustrate its key implications in a stylized two-period model of firm investment and borrowing.

2.1 Diagnostic Expectations

Starting from the 1970s, psychologists Kahneman and Tversky (KT) assembled extensive laboratory evidence that human judgments depart in systematic ways from Bayesian updating. For instance, individuals tend to neglect base rates, leading them to over-react to noisy signals. KT accounted for this and other biases by proposing that beliefs are often formed using "representative" information that easily comes to mind.[3] To give one example, Casscells et al. (1978) showed that physicians tend to exaggerate the probability that a patient is sick after a positive medical test. In KT's logic, this mistake is due to the fact that after a positive test the "sick" patient type

[2]This IV specification is simply a useful tool for organizing the timing of variation in the data, and our goal here is not the achievement of reduced-form causal identification per se. We later build a model in which the endogenous patterns we present here can be explained with diagnostic beliefs but not with rational beliefs. In other words, these patterns are useful for structural identification even though the variables of interest are clearly endogenous.

[3]Kahneman and Tversky also discuss other heuristics that can lead to distorted judgments, such as "availability" and "anchoring". See Gennaioli and Shleifer (2018) for a systematic discussion of this work.

easily comes to mind, interfering with recall of healthy types. Thus, the physician neglects the fact that the disease is rare, causing an inflated assessment of sickness.

Tversky and Kahnemann offer the following definition of representativeness (Tversky and Kahneman, 1983): "an attribute is representative of a class if it is very diagnostic; that is, if the relative frequency of this attribute is much higher in that class than in the relevant reference class."

Here a class could be a social group we seek to judge, a set of firms we seek to assess, or the current state of the economy on the basis of which we predict the future. Its representative traits are then those that are relatively more frequent in this class than in other classes. For instance, when thinking about the Irish, red hair is representative because it is relatively more frequent in this class than in the other national groups.

Building on this definition, we offer a model of beliefs in which representative traits are more easily recalled and thus overweighed in judgments. In this case, beliefs overreact in the direction of events that have become relatively more likely, even if they are rare in absolute terms. For instance, we inflate the probability of sickness because the positive medical test increases its objective probability *relative* to not observing any test outcome. Critically, in this approach belief distortions depend on the true statistical features of the data generating process. This property, which psychologists call "the kernel of truth," allows for disciplined applications of the model.[4]

When applied to dynamic contexts, the model works as follows. Suppose that a variable follows an AR(1) process $X_{t+1} = \rho X + \varepsilon_{t+1}$ where ε_{t+1} is Gaussian with mean zero and standard deviation σ. At time t the agent seeks to forecast X_{t+1}. The representativeness of a realization X_{t+1} is assumed to be measured by the likelihood ratio:

$$R_t\left(X_{t+1}\right) = \frac{f\left(X_{t+1}\left|X_t\right.\right)}{f\left(X_{t+1}\left|\rho X_{t-1}\right.\right)}, \tag{2}$$

where $f\left(X_{t+1}\left|.\right.\right)$ denotes the density of X_{t+1} conditional on a value at t . As in the KT definition, the future realization X_{t+1} is highly representative when its probability increases a lot on the basis of recent news $\varepsilon_t = X_t - \rho X_{t-1}$.[5]

The diagnostic distribution of X_{t+1} is then defined as:

$$f^{\theta}\left(X_{t+1}\left|X_t\right.\right) \propto f\left(X_{t+1}\left|X_t\right.\right)\left[\frac{f\left(X_{t+1}\left|X_t\right.\right)}{f\left(X_{t+1}\left|\rho X_{t-1}\right.\right)}\right]^{\theta}, \tag{3}$$

[4]In previous work we showed that this organizing principle offers a parsimonious account of measured belief distortions in diverse domains such as social stereotypes (Bordalo et al., 2016, 2019a), long-term earnings growth estimates (Bordalo et al., 2018, 2019b), and macroeconomic forecasts (Bordalo et al., Forthcomingb).

[5] In Equation (2) the reference distribution in the denominator conditions on the absence of the recent shock ε_t. As we discuss in Bordalo et al. (2018), the reference distribution could be formalized as the density of X_{t+h} conditional on yesterday's observation X_{t-1}, or it could be defined in terms of a weighted average of past realizations, reflecting the influence of more remote memories. Notwithstanding the differences, the intuition is the same: good recent news increase the extent to which good future outcomes are representative. The current specification is very tractable and displays some convenient formal properties, which is why we use it here. We leave the exploration of alternative specifications to future work.

where $\theta \geq 0$ parameterizes the extent of belief distortions. The model nests rational expectations as the special case in which $\theta = 0$. For $\theta > 0$ agents inflate the probability of future outcomes that have become more likely in light of recent data.

Equation (3) should be interpreted as the product of selective memory. The agent's database contains in principle all necessary information to make a correct judgment, summarized by the true conditional distribution $f(X_{t+1}|X_t)$. However, more representative information more easily comes to mind, causing the agent to inflate its weight in judgments. This effect is captured by the likelihood ratio in (3). Because the database contains objectively useful information, the beliefs in Equation (3) are forward looking: they depend entirely on the true data generating process. As a result, expectations react to regime changes, circumventing the Lucas (1976) critique of adaptive expectations. As we will see, this also implies that belief updating depends on intrinsic characteristics of the series, such as its persistence ρ.

When the true conditional distribution is Gaussian, as assumed above, it can be shown (Bordalo et al., 2018) that the diagnostic distribution $f^{\theta}(X_{t+1}|X_t)$ is also Gaussian, with the same standard deviation σ of the true distribution, and the distorted mean:

$$\mathbb{E}_t^{\theta}(X_{t+1}) = \mathbb{E}_t(X_{t+1}) + \theta\left[\mathbb{E}_t(X_{t+1}) - \mathbb{E}_{t-1}(X_{t+1})\right], \tag{4}$$

where $\mathbb{E}_t(.)$ denotes the rational expectation. The diagnostic expectation of X_{t+1} can be interpreted as a distortion of the rational expectation $\mathbb{E}_t(X_{t+1})$, which corresponds to an unbiased usage of the memory database, toward states that have become more likely in light of news. When news are positive, the agent is excessively optimistic. When news are negative, he is excessively pessimistic. On average, news are zero, so expectations fluctuate around the rational benchmark.

In the current AR(1) setting, this implies that at time t the diagnostic distribution of X_{t+1} is normal with variance σ^2 and mean $\rho X_t + \rho\theta\varepsilon_t$, as if at time t the agent believes that the state follows the ARMA (1,1) process:

$$X_{t+1} = \rho X_t + \rho\theta\varepsilon_t + \varepsilon_{t+1}. \tag{5}$$

In this sense, Diagnostic Expectations generate a form of model misspecification. This misspecification is not mechanical. By the kernel of truth logic, the believed model depends on the true process, which here is an AR(1), and in particular on its true persistence ρ and volatility σ^2. By linking beliefs to measurable features of reality, diagnostic expectations offer a disciplined approach to analyzing departures from rational expectations.

Finally, although diagnostic expectations introduce misspecification, our focus on the dynamics of beliefs themselves separates our approach from the distinct but complementary literature in macroeconomics on model misspecification and robust decisionmaking in such contexts. See, for example, the survey in Hansen and Sargent (2001).

2.2 A two period model of risky borrowing and investment

To see how diagnosticity intuitively affects borrowing and investment, consider a stylized two period model. A risk neutral and patient entrepreneur transforms capital into output using technology $y = Ak^{\alpha}$, with $\alpha < 1$. Capital is installed one period in advance, and future productivity A is lognormally distributed, $\ln A = \rho \ln A_0 + \varepsilon$, where A_0 is current productivity. To invest, the entrepreneur issues one period bonds, each promising a unit face value and fetching a market price q . The entrepreneur borrows just enough to finance investment, $k = qb$, where b is the amount of bonds issued and the future repayment.

Debt can default. If the entrepreneur does not repay, output is lost in deadweight bankruptcy costs and all parties recover zero. If the entrepreneur repays, he keeps the firm's profit. It is then optimal for the entrepreneur to default if output is less than debt b plus any added default penalty.[6] We consider two extreme cases: one in which default penalties are infinite, so that debt is riskless, and another in which they are absent.

Lenders are deep pocketed and risk neutral: they are willing to supply any amount of funds provided their perceived expected return is equal to the riskless interest rate, which we assume to be zero. Thus the price of debt is $q = \pi$, where π is the probability of future repayment.

2.2.1 Riskless Debt

With infinite default penalties repayment is certain, $\pi = 1$, so the price of debt is one, $q = 1$. The entrepreneur chooses how much to borrow and invest (recall that $k = b$) to maximize expected output minus investment costs:

$$\max_{b} -b + \mathbb{E}^{\theta}(A)\, b^{\alpha}.$$

Exploiting the lognromality of TFP and Equation (5), we can show that the log deviation of optimal investment and borrowing from to the rational benchmark b^* is given by:

$$\ln\left(\frac{b}{b^*}\right) = \left(\frac{1}{1-\alpha}\right)\theta\rho\varepsilon. \tag{6}$$

After good news, $\varepsilon > 0$, the entrepreneur is too optimistic, particularly if productivity is very persistent (high ρ). He borrows and invests more than in the rational expectations case. After bad news, $\varepsilon < 0$, the reverse occurs. Diagnosticity amplifies booms and busts. Because borrowing and investment are shaped not only by current productivity but also by the recent productivity path, the contemporaneous correlation between productivity and macro outcomes is also reduced relative to the RE case.

[6] Debt repayment can be larger than output because at $t+1$ the entrepreneur receives a large abscondable endowment.

A third, critical, implication of diagnosticity is that it creates systematic reversals that are predictable to the (non diagnostic) econometrician. In the current two period framework, reversals can be obtained by averaging Equation (6) over the true distribution of news:

$$\mathbb{E}\left[\ln\left(\frac{b}{b^*}\right)\right] = 0 \ .$$

Because on average there are no news, diagnostic expectations are on average rational. As a consequence, periods in which borrowing and investment are excessively high are systematically followed by downward corrections and vice versa.

Finally, due to initial over-reaction and systematic reversals, diagnostic expectations cause leverage and investment to be excessively volatile, and the more so the higher is θ :

$$Var\left[\ln\left(\frac{b}{b^*}\right)\right] = \left(\frac{1}{1-\alpha}\right)^2 \theta^2\rho^2\sigma^2,$$

where again the variance is computed using the true distribution of TFP shocks.

The magnitude of these effects depends on the economic environment. Diagnosticity exerts stronger effects when: i) uncertainty σ^2 is higher, ii) returns to scale are slowly diminishing ($\alpha \to 1$), and iii) TFP is more persistent (ρ is higher). When we later quantify the macroeconomic effects of diagnosticity, these factors make a significant difference.

2.2.2 Risky Debt

Without default penalties, there are states in which the entrepreneur optimally defaults, so debt is risky. Now diagnosticity has two additional implications. First, it affects lenders' expectations about repayment and thus the equilibrium price of debt. Second, it changes the entrepreneur's perception of default risk, further affecting borrowing and investment.

Consider lenders' expectations about debt repayment. The entrepreneur repays when output is higher than debt, $y > b$, which occurs when productivity is high enough, $A > \pi^{-\alpha}b^{1-\alpha}$. With the diagnostic lognormal distribution of productivity entailed by (5), the perceived probability of repayment π and hence the price of debt $q = \pi$ is implicitly defined by the equation:

$$\pi = 1 - \Phi\left(\frac{-\alpha\ln\pi + (1-\alpha)\ln b - \mathbb{E}^{\theta}(\ln A)}{\sigma}\right), \tag{7}$$

where $\Phi(.)$ is the Gaussian cdf.

As we show in the Appendix, if debt b is sufficiently low, there is a unique and interior *perceived* probability of repayment $\pi(b)$ that decreases in b . Intuitively, the higher the level of debt, the lower the probability of repayment.[7]

[7]Equation (7) also allows for a "zero debt" equilibrium in which for any positive debt level default occurs with probability one, $\pi = 0$. We rule out this equilibrium because it is unstable.

The perceived probability of repayment in (7) depends also on beliefs about productivity $\mathbb{E}^{\theta}(\ln A)$. When news is good, lenders' are too optimistic and the perceived probability of repayment π is too high. Default risk is neglected at any level of debt b and the price of debt is then too high. Conversely, during bad times risk perceptions are heightened and the price of debt is too low. In a cross section of firms that vary according to their idiosyncratic productivities and hence their default risk, this mechanism pins down credit spreads.

The mispricing of debt in turn feeds into entrepreneurs' borrowing and investment decisions. With risky debt, the entrepreneur's problem becomes:

$$\max_{b} -\pi(b)\, b + \pi(b)^{\alpha}\, b^{\alpha} \int_{\pi(b)^{-\alpha} b^{1-\alpha}}^{+\infty} A f^{\theta}(A \,|A_0)\, dA. \tag{8}$$

The first term captures investment costs, which are equal to the amount of funds raised. The second term captures expected output, whose distribution is truncated in the event of default. There are two noteworthy features. First, the entrepreneur himself internalizes default risk: he wishes to reduce the probability that the firm is shut down. This is the standard bankruptcy cost of debt, which induces the entrepreneur to restrain leverage and investment. Second, the entrepreneur also internalizes the fact that by issuing more debt b he reduces the perceived probability of repayment, thereby reducing the price of debt $\pi(b)$. Effectively, the entrepreneur faces a debt Laffer curve $\pi(b)\, b$.

Diagnosticity shapes the market calculus in two ways. First, it distorts the entrepreneur's perception of future productivity, distorting his incentive to invest, and hence his demand for funds. This effect is captured by the rightmost integral in (8). Second, it distorts lenders' perception of default and risk and hence the price schedule $\pi(b)$. For instance, when lenders are too optimistic, the price of debt is too high, which induces the entrepreneur to issue and invest even more. In this sense, diagnosticity affects debt and investment by shaping both the demand and the supply side of capital. One important question here is which one contributes more. We investigate this issue in our model.

3 An RBC Model with Diagnostic Firms and Lenders

We now introduce our quantitative business cycle model with TFP driven fluctuations. We begin with a partial equilibrium setup. Firms with different and persistent productivities decide whether to default, to hire labor, to borrow, and to invest subject to capital adjustment costs. Credit is supplied by a continuum of risk neutral lenders. The riskless rate R and the wage rate W are given. The only difference relative to a workhorse model with firm heterogeneity and risky debt (Khan and Thomas, 2008; Arellano et al., 2019; Gilchrist et al., 2014), is that firms and lenders form expectations diagnostically.

Time is discrete. We use $'$ to denote future values and $_{-1}$ to indicate lagged values. Uppercase letters refer to aggregate or common values, lowercase letters refer to idiosyncratic objects.

3.1 Firms

The generic firm has micro-level TFP z and is subject to macro level productivity A. It uses capital k and labor n as inputs to produce output according to a decreasing returns technology

$$y = Azk^{\alpha}n^{\nu}, \qquad \alpha + \nu < 1.$$

The log of micro TFP follows the AR(1) process

$$\log z' = \rho_z \log z + \varepsilon'_z, \quad \varepsilon'_z \sim N(0, \sigma_z^2), \quad 0 < \rho_z < 1 \ , \tag{9}$$

while the log of macro TFP follows a similar process with

$$\log A' = \rho_A \log A + \varepsilon'_A, \quad \varepsilon'_A \sim N(0, \sigma_A^2), \quad 0 < \rho_A < 1 \ . \tag{10}$$

Firms invest i in capital k with one-period time to build

$$k' = i + (1 - \delta)k, \qquad 0 < \delta < 1.$$

Investment requires quadratic adjustment costs $AC(i, k) = \frac{\eta}{2}\left(\frac{i}{k}\right)^2 k$ indexed by parameter $\eta > 0$.

Firms act competitively. In each period, the timing of events is as follows. First, each firm decides whether to default on its debt. If a firm defaults, its assets are transferred to lenders. If a firm repays, it hires labor at the wage W, and chooses how much to invest and how much new one period debt to issue. It does so to maximize the discounted sum of current and future dividends, where the discount rate $(1 + R)^{-1} < 1$ reflects the exogenous risk-free rate R.

The firm's current dividend d must be non-negative (we rule out equity issuance) and is given by:

$$d = (1 - \tau)\left[y - Wn - AC(i, k) - \phi\right] + q^{\theta}(s, k', b')b' - i - b + \tau(R + \delta k). \tag{11}$$

The firm's earnings are given by its output minus the wage bill, the adjustment cost, and a fixed production cost $\phi > 0$, and it pays a corporate income tax rate $\tau \in (0, 1)$ on them. The firm raises additional resources by issuing new debt b', priced by the schedule $q^{\theta}(s, k', b')$, it incurs the investment cost i , and repays its current debt b . Finally, the firm receives tax rebates for capital depreciation and interest expenses on debt.[8] This formulation of dividends and specification of firm fundamentals is standard, e.g. Strebulaev and Whited (2012).

To decide whether to default and how much to borrow and invest, the firm forms beliefs about its future productivity. To assess default risk and interest rates, lenders must do the same.

[8]For computational simplicity, we assume the rebate is on average equal to the cost of debt R.

We assume that both firms and lenders form expectations diagnostically. Given the true AR(1) processes (9) and (10), and given the diagnostic formula in (5), diagnostic beliefs over micro and macro TFP are described by the normal processes:

$$\log z'|(\log z, \varepsilon_z) \sim N\left[\rho_z(\log z + \theta\varepsilon_z), \sigma_z^2\right] \tag{12}$$

$$\log A'|(\log A, \varepsilon_A) \sim N\left[\rho_A(\log A + \theta\varepsilon_A), \sigma_A^2\right]. \tag{13}$$

With diagnostic expectations, $\theta > 0$, the agent forecasts future productivity by overweighting current news, as if the true productivity process follows an ARMA (1,1).[9]

In this case, when forming beliefs about a generic firm, diagnostic agents consider four state variables: its current micro TFP z , aggregate macro TFP A , the micro shock ε_z and the macro shock ε_A . We collect these exogenous states in the vector $s = (z, \varepsilon_z, A, \varepsilon_A)$. A firm is also identified by two endogenous states, its inherited capital stock k and debt b. Given an overall state (s, k, b), the firm defaults if its diagnostically expected continuation value is negative, and it repays otherwise. If the firm repays, it hires labor, invests, and borrows so as to maximize the sum of the current and diagnostically expected discounted future earnings, taking into account the possibility of default in the future.

This problem can be written in a recursive fashion. Upon entering the current period, the value of the firm is given by:

$$V^\theta(s, k, b) = \max\left[0, V_{ND}^\theta(s, k, b)\right], \tag{14}$$

where $V_{ND}^\theta(s, k, b)$ is the continuation value from non defaulting. Condition $V^\theta(s, k, b) = 0$ identifies states in which the firm optimally defaults. The continuation value from non defaulting is recursively determined as:

$$V_{ND}^\theta(s, k, b) = \max_{k', b', n, \text{ s.t. } d \geq 0} d + \frac{1}{1+R}\mathbb{E}^\theta\left[V^\theta(s', k', b')|s\right]. \tag{15}$$

If the firm does not default, it optimally hires labor n , sets future capital k' and debt b' so as to maximize its current dividend plus its diagnostically expected discounted future value $V^\theta(s', k', b')$.

Relative to the rational expectations benchmark, diagnosticity introduces two modifications. First, the debt price $q^\theta(s, k', b')$ is determined by competitive lenders whose expectations are diagnostic. Second, the firm considers a diagnostic expectation of its future value, captured by the notation $\mathbb{E}^\theta(.)$.[10]

[9] Another approach to capture extrapolation is Fuster et al. (2010)'s Natural Expectations, in which long lags in the data generating process are neglected by agents who end up overestimating short-term persistence in processes with long-term mean reversion. In the current AR(1) setting, such beliefs would be indistinguishable from RE.

[10] We apply diagnostic expectations to the recursive formulation of the problem, Equation (15). The diagnostic

The labor choice n is statically optimized out of a non-defaulting firm's dynamic decision problem, leaving only the intertemporal choices of k' and b'. More generally, Equations (14) and (15) determine both the optimal firm default policy by $df^\theta(s,k,b)$ and the policies for endogenous states $k'^\theta(s,k,b)$, $b'^\theta(s,k,b)$.

3.2 Lenders

Firms borrow from risk-neutral deep-pocket lenders who form expectations diagnostically and require an expected return equal to the risk-free rate R. If a firm (s,k,b) defaults on its debt b, the lender receives the recovery rate

$$\mathcal{R}(s,k,b) = (1-\tau)\,\gamma \frac{\max\left[y - Wn - \phi, 0\right] + (1-\delta)k}{b}$$

which reflects, net of tax, an exogenous fraction γ of the firm's cash flows and the liquidation value $\gamma\,(1-\delta)\,k$ of its remaining capital stock. The remaining fraction $1-\gamma$ is a deadweight loss. In the event of default, a reorganized firm carries on into the next period with zero capital and debt, in which case it is forced to again borrow to finance investment and future growth.

The price of debt $q^\theta(s,k',b')$ adjusts endogenously so that the diagnostically expected bond return is equal to the risk free rate R . Formally this means that:

$$q^\theta(s,k',b') = \frac{1}{1+R}\mathbb{E}^\theta\left[1 + df^\theta(s',k',b')\left(\mathcal{R}(s',k',b') - 1\right)|s\right]. \tag{16}$$

To equalize expected bond return across firms, riskier firms must promise a higher interest rate. By this logic, the firm's interest rate spread relative to the risk-free rate is given by:

$$S^\theta(s,k',b') = \frac{1}{q^\theta(s,k',b')} - (1+R).$$

These equations illustrate how diagnosticity affects debt prices and spreads. On the demand side for capital, it affects the firm's default $df^\theta(s,k,b)$, debt $b'^\theta(s,k,b)$ and investment $k^\theta(s,k,b)$ policies. On the supply side, it affects the probability of default perceived by lenders, as captured by the operator $\mathbb{E}^\theta$ in (16). We later analyze how these demand and supply forces contribute to the credit cycle.

3.3 Solving the Model

A solution to the model reflects a set of firm level policies and values $b'^\theta, k'^\theta, df^\theta, V^\theta_{ND}, V^\theta$ together with a set of debt price schedules q^θ. These objects must jointly satisfy optimization by firms, Equations (14) and (15), as well as the lenders' zero-profit condition in Equation (16).

agent believes that productivity follows an ARMA (1,1) and correctly thinks that he will continue to believe the same in the future. The recursive problem is equivalent to an optimal control problem in which the probability distribution of A_{t+s} at time t is the product $\Pi_{j=1}^{s} f^\theta(A_{t+j}|A_{t+j-1},\epsilon_{t+j-1})$ of the conditional distributions between times t and $t+s-1$. This distribution has the same mean as the time t diagnostic distribution $f^\theta(A_{t+s}|A_t,\epsilon_t)$ but has larger variance. This is due to overreaction to news (which are zero on average) in the intermediate periods.

We solve the model numerically. In addition to standard Bellman equations V^θ and V^θ_{ND}, the model features a fixed point between firm default policies df^θ and credit prices q^θ. To solve the model, we employ an iterative approach detailed in Appendix A. First, we guess a firm default rule df^θ, computing the implied debt price schedule q^θ according to the lenders' zero-profit condition above. Then, we compute the solution to a firm's dynamic problem by solving the Bellman equations V^θ and V^θ_{ND} using discretization and policy iteration. If the implied default states, i.e., those states with negative value $V^\theta < 0$, match the set of initial guesses, the iteration is complete. Otherwise, we compute the newly implied default states and repeat the process. The algorithm we employ is standard within the literature solving quantitative dynamic corporate finance models and follows the implementation in Strebulaev and Whited (2012).[11]

To illustrate how diagnosticity affects firm level choices, Figure 2 plots the value function V^θ in our baseline calibration (which we discuss below), as a function of capital k. Each line in the figure plots the perceived value for a firm with different realizations of idiosyncratic news ε_z but otherwise identical states $(z, A, \varepsilon_A, k, b)$.

The message echoes the two period model of Section 2. After good news, diagnostic firms are too optimistic, so they value capital more than an otherwise identical firm with ratinoal expectations. These overoptimistic firms invest and borrow more than their rational counterparts. This mechanism proves crucial for generating macro level effects: after good aggregate news, overinvestment and over-borrowing by diagnostic firms render the economy vulnerable to crises. The reverse is true for bad aggregate shocks, which create excess pessimism, deep crises, but also predictable recoveries.

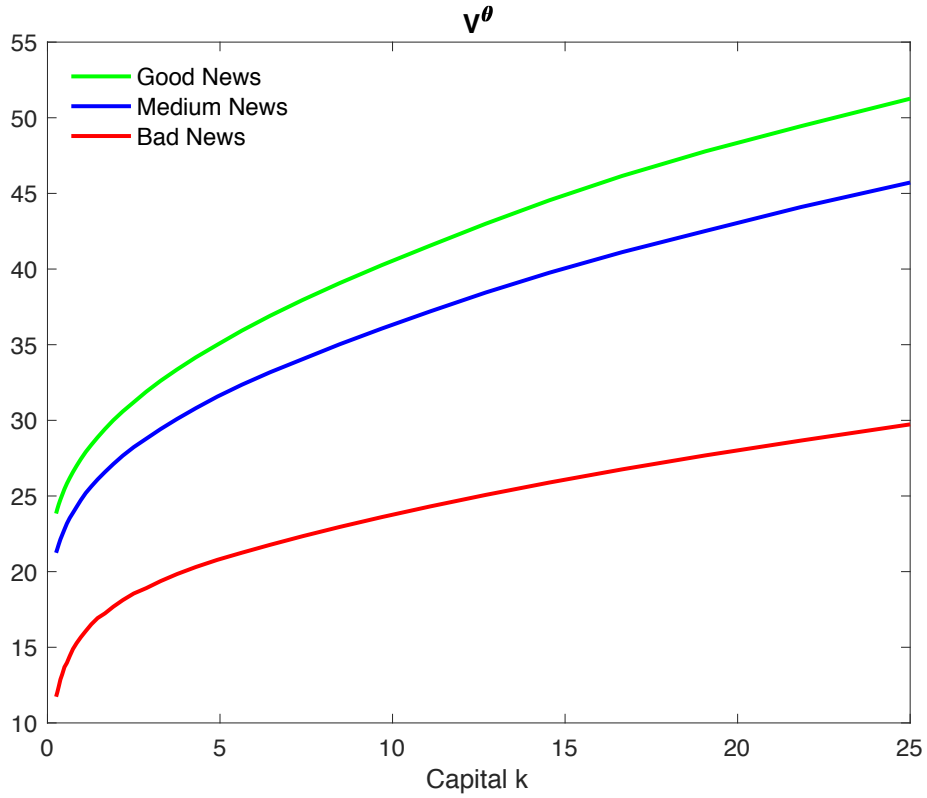
4 Model Calibration

We set model parameters in two steps. First, we match eight parameters to conventional values for a model like ours solved at an annual frequency. Table 3 reports the eight externally fixed parameters. Given the similarity of the production structure and macro TFP fluctuations, we draw on Bloom et al. (2018) for a range of firm-level and macro TFP parameters. Information on effective corporate income taxes is obtained from the Congressional Budget Office (CBO, 2017).

Second, we calibrate the remaining six parameters by matching moments computed from micro-data. These parameters govern the micro-level TFP process ρ_z and σ_z, adjustment and operating costs η and ϕ, lender recovery rates γ, and the diagnosticity parameter θ. We set these parameters to best match thirteen moments on firm level investment rates, profits, leverage, default

[11] Our numerical approach here is highly computationally intensive, given the presence of four exogenous states, two endogenous states/policies, and endogenous default rule, and a pricing fixed point. However, judicious application of parallelization and an economical approach to storage of micro-level outcomes following Young (2010) and Terry (2017a) allow for solution of the model in several minutes in a standard university cluster computing environment using Fortran.

Figure 2: Firm Value Exhibits Diagnosticity



Notes: The figure plots the perceived diagnostic value function V^{θ} for firms as a function of capital k for the calibrated DE model. The three lines correspond to identical current productivity levels z but distinct values of news (green line after good news, blue line after moderate news, and red line after negative news).

rates, credit spreads, and crucially on the predictability of errors in firm forecasts of their own earnings. We note at the outset that this is a highly overidentified moment-matching exercise with a nonlinear model. It allows us to exploit information on many moments, but of course we are not in general able to deliver an exact fit.

We also note that while many of these parameters and moments are familiar in the corporate finance and macroeconomics literatures, the moments involving forecast errors and their link to the belief parameter θ deserve further discussion. In the RE model with $\theta = 0$, future forecast errors should be unpredictable using any currently available information. In contrast, in the DE model with $\theta > 0$, overreaction of expectations towards the direction of recent news leads to systematic reversals. When times are good today, managers are likely to be disappointed tomorrow, and vice-versa, driving a negative correlation between future forecast errors and the level of today's profits. This makes data on the sign and magnitude of forecast errors a good way to pin down the value of θ.

Table 4 reports the values of the thirteen targeted moments. We obtain annual financial

Table 3: Externally Fixed Parameters

	Parameter	Value	Explanation	Source
1	δ	0.1	Depreciation rate	Annual solution
2	R	0.04	Risk-free rate	Annual solution
3	α	0.25	Capital revenue elasticity	Bloom et al. (2018)
4	ν	0.50	Labor revenue elasticity	Bloom et al. (2018)
5	ρ_A	0.9	Macro TFP persistence	Bloom et al. (2018)
6	σ_A	0.015	Macro TFP shock standard deviation	Bloom et al. (2018)
7	W	0.5	Wage	Normalization
8	τ	0.20	Corporate income tax	Effective corporate tax rates, CBO (2017)

Notes: The table reports the parameter symbol, numerical value, a description, and source information for each of the externally fixed parameters. Outside of the unit-free persistence or normalized parameters, all reported values are in proportional units, e.g. $0.01 = 1\%$.

statements for listed US firms from Compustat, extracting information on earnings, investment, debt, and capital. The exact definitions, variables used, and details of sample construction are available in Appendix B.

With this standard data in hand, we construct a dataset of profit forecasts made by the managers of the same firms from the IBES guidance database considered above in Section 1. With these forecasts in hand, we can construct measures of profit forecast errors equal to realized profits minus the manager's forecast made one period ago. Once again, further details on the construction of our dataset are available in Appendix B, together with descriptive statistics of our sample of firms.

After merging the Compustat and IBES samples, the resulting panel spans 2007-2016 for about 1000 firms and approximately 5000 observations, with a sample span reflecting a time period with particularly large numbers of firms reporting manager expectations and a selection of mostly large firms closely watched by financial markets and investors. We normalize firm profits π, investment i, issued debt b', and forecast errors in the next period fe' by firm capital k, and we target the covariance matrix of all of these series. The 11th moment is the autocorrelation of profit rates in the same data. Finally, we draw average corporate default rates from Gourio (2013) and interest rate spreads over risk-free debt from Moody's BAA-treasury series. In line with the empirical evidence in Hamilton and Cantor (2005), we define the BAA credit spread in the model as the median firm's credit spread. Data Appendix B contains an extensive discussion of the samples and definitions used empirically and in the model simulations.

We choose this set of moments for several reasons. Firm scaled earnings $\frac{\pi}{k}$ and their correlations encode information about the productivity process at firms, helping to identify σ_z and ρ_z. Firm investment rates $\frac{i}{k}$ not only reflect the productivity processes but also the various frictions such as adjustment costs, helping to identify η. Leverage choices $\frac{b'}{k}$ reveal crucial information about

Table 4: Target Moments

	Moment	Value	Explanation	Source
1	$\sigma\left(\frac{\pi}{k}\right)$	0.3406	Standard Deviation of Profit	Compustat-IBES sample
2	$\rho\left(\frac{\pi}{k}, \frac{i}{k}\right)$	0.3128	Correlation of Profit, Investment	Compustat-IBES sample
3	$\rho\left(\frac{\pi}{k}, \frac{b'}{k}\right)$	0.1518	Correlation of Profit, Leverage	Compustat-IBES sample
4	$\rho\left(\frac{\pi}{k}, \frac{fe'}{k}\right)$	-0.2141	Correlation of Profit, Forecast Error	Compustat-IBES sample
5	$\sigma\left(\frac{i}{k}\right)$	0.1124	Standard Deviation of Investment	Compustat-IBES sample
6	$\rho\left(\frac{i}{k}, \frac{b'}{k}\right)$	0.1791	Correlation of Investment, Leverage	Compustat-IBES sample
7	$\rho\left(\frac{i}{k}, \frac{fe'}{k}\right)$	-0.1237	Correlation of Investment, Forecast Error	Compustat-IBES sample
8	$\sigma\left(\frac{b'}{k}\right)$	0.5346	Standard Deviation of Leverage	Compustat-IBES sample
9	$\rho\left(\frac{b'}{k}, \frac{fe'}{k}\right)$	-0.0594	Correlation of Leverage, Forecast Error	Compustat-IBES sample
10	$\sigma\left(\frac{fe'}{k}\right)$	0.2347	Standard Deviation of Forecast Error	Compustat-IBES sample
11	$\rho\left(\frac{\pi}{k}, \frac{\pi_{-1}}{k_{-1}}\right)$	0.6637	Autocorrelation of Profit	Compustat-IBES sample
12	$\mathbb{E}$Default	0.0050	Mean Default Rate	Gourio (2013)
13	$\mathbb{E}$Spread	0.0320	Mean Interest Rate Spread	Moody's

Notes: The table reports the target moments used in the calibration of the model. The first 11 moments are drawn from a sample of US listed firms combining data from Compustat and IBES at the firm-fiscal year level spanning 2007-2016 for 867 firms and 4457 observations. σ refers to standard deviations, while ρ refers to correlations. The normalizer for all series k is tangible capital (the book value of plants, property, & equipment). The profit series π is Street or pro forma earnings. The investment series i is capital expenditures. Debt b' is total liabilities at the end of period. The forecast error fe' is the next-fiscal year value of pro-forma earnings minus manager forecasts of earnings. The default moment is drawn from Gourio (2013). The mean spread is the average value of the Moody's BAA-Treasury annualized spread. Outside of the unit-free correlations, all reported values are in proportional units, e.g. $0.01 = 1\%$.

the beliefs of firms and of investors, aiding in the identification of θ. Future forecast errors $\frac{fe'}{k}$ are especially revealing of systematic errors in beliefs at the micro level, mapping naturally to the diagnosticity parameter θ as noted above. Finally, mean default and spread values encode information about both the fixed costs of operation ϕ and recovery fractions γ conditional upon default. To calibrate the six model parameters, we minimize the deviation between the empirical moments in Table 4 and those computed from a comparable unconditional simulation of the model. We weight the moments by one over their squared value, implying an objective function equal to the sum of squared percentage deviations between model and data moments.

One key feature of the microdata is strong evidence of forecast error predictability, as reflected in the negative correlation between current firm profits $\frac{\pi}{k}$ and future forecast errors $\frac{fe'}{k}$ in row 4 of Table 4. Recall that these forecast errors fe' are the difference between next year's realized profits and current forecasts of those profits. This means that good current conditions, as measured by high profits $\frac{\pi}{k}$, predict systematic future disappointment as measured by negative forecast errors $\frac{fe'}{k}$. This correlation is quite robust, with a t-statistic of -7.7. Intuitively, given their

positive correlation with today's profits, higher investment and debt issuance today also predict future disappointment or negative forecast errors in rows 7 and 9 of Table 4.[12] Similar predictable reversions relative to expectations are documented in Gennaioli et al. (2016) for firm level earnings and in Bordalo et al. (Forthcomingb) for a range of macro forecasts.[13] This evidence is at odds with rational expectations, and we now use the model structure to determine whether the magnitude of these predictable reversals imply meaningfully sized values of $\theta > 0$.

4.1 Calibrated Parameters

Table 5: Internally Calibrated Parameters

	Parameter	Value	Explanation
1	θ	1.076	Diagnosticity parameter
2	ρ_z	0.882	Micro TFP persistence
3	σ_z	0.103	Micro TFP shock standard deviation
4	η	2.839	Capital adjustment cost
5	ϕ	0.130	Fixed operating cost
6	γ	0.164	Recovery rate

Notes: The table reports the parameter symbol, calibrated value, and an intuitive explanation for each of the internally calibrated parameters. These 6 parameters were fixed by targeting the values of the 13 empirical moments. During the calibration process, the model equivalents of the empirical moments were drawn from a simulated sample of 2500 firms over 10 years, approximately twice the size of the empirical sample. The simulation was performed using a set of unconditional draws of shocks for the model held constant across parameter values, and moments were computed after an initial 250-year simulation to reduce influence of initial conditions on the resulting moments.

Table 5 reports the parameters for the calibrated DE model. We later discuss how these values fit the data moments. The diagnosticity parameter $\theta \approx 1$ matches closely the values found by Bordalo et al. (2018) using data on professional forecasts of credit spreads ($\theta = 0.9$), by Bordalo et al. (2019b) using analyst expectations of US listed firms' long term earnings growth ($\theta = 0.9$), and by Bordalo et al. (Forthcomingb) using professional forecasts of several macro series ($\theta = 0.6$). A value of θ close to 1 means that forecast errors are roughly equal to the size of incoming news. The calibrated values of micro TFP persistence ρ_z and volatility σ_z are close to those from other work calibrating or structurally estimating firm-level shock processes with Compustat data, e.g., Hennessy and Whited (2007), Gourio and Rudanko (2014), Terry (2017b), Khan and Thomas (2008), or Saporta-Eksten and Terry (2018).

[12] The t-statistic quoted above for row 4 reflects standard errors clustered by firm. The firm-clustered t-statistics for rows 7 and 9 are -6.02 and -3.75, respectively.

[13] Given our focus on manager decision-making, we study manager beliefs using earnings guidance data from IBES. Bouchaud et al. (2019) study equity analysts' short-term earnings forecasts which, while correlated with the former, display a form of underreaction. We note that both the variable forecasted and the incentives involved are distinct in the latter sample.

4.2 Fit with Microdata

Table 6: Model vs Data Moments

	Moment	Data	Diagnostic (DE) Model	Rational (RE) Model
1	$\sigma\left(\frac{\pi}{k}\right)$	0.3406	0.0962	0.1080
2	$\rho\left(\frac{\pi}{k}, \frac{i}{k}\right)$	0.3128	0.1088	0.5759
3	$\rho\left(\frac{\pi}{k}, \frac{b'}{k}\right)$	0.1518	0.4819	0.8235
4	$\rho\left(\frac{\pi}{k}, \frac{fe'}{k}\right)$	-0.2141	-0.1820	-0.0014
5	$\sigma\left(\frac{i}{k}\right)$	0.1124	0.1099	0.0873
6	$\rho\left(\frac{i}{k}, \frac{b'}{k}\right)$	0.1791	0.6035	0.4267
7	$\rho\left(\frac{i}{k}, \frac{fe'}{k}\right)$	-0.1237	-0.2459	-0.0061
8	$\sigma\left(\frac{b'}{k}\right)$	0.5346	0.2196	0.2397
9	$\rho\left(\frac{b'}{k}, \frac{fe'}{k}\right)$	-0.0594	-0.2483	-0.0068
10	$\sigma\left(\frac{fe'}{k}\right)$	0.2347	0.0746	0.0564
11	$\rho\left(\frac{\pi}{k}, \frac{\pi'}{k'}\right)$	0.6572	0.7215	0.8533
12	$\mathbb{E}$Default	0.0050	0.0046	0.0065
13	$\mathbb{E}$Spread	0.0320	0.0568	0.0243

Notes: The first column defines the moments. σ is standard deviation, ρ is correlation. The second column reports empirical moments. The third column reports moments from the calibrated DE model. The fourth column reports moments from the RE model with $\theta = 0$. All reported values are in proportional units, e.g. $0.01 = 1\%$. The first 11 moments are from the microdata Compustat-IBES sample spanning 2007-2016. The spread is the Moody's BAA-Treasury spread, and the default rate is from Gourio (2013). Simulated moments from 2500 firms over 50 years.

Table 6 reports the fit of the micro moments for both the calibrated DE model and an associated RE model with $\theta = 0$ but all other parameters unchanged. Comparison across the two models isolates the impact of the diagnosticity parameter θ on simulated behavior. As usual in highly overidentified exercises with nonlinear models such as this one, the calibrated DE model does not fit perfectly. However, the fit is comparable to similar exercises in structural corporate finance or quantitative macroeconomics (Hennessy and Whited, 2007; Terry, 2017b; Bloom et al., 2018). Crucially, the DE model also improves upon the RE model along several dimensions.

Most intuitively, $\theta > 0$ allows to match the predictability of errors. Diagnostic over-reaction to news offers a good fit for the evident excess managerial optimism (pessimism) in good (bad) times. The RE model cannot account for these predictable errors, and in fact when $\theta = 0$ these correlation are not meaningfully different from zero.

The DE model is also better able to capture the correlation of current profitability with investment and debt, and the autocorrelation of profitability. As we saw in Section 2, diagnostic investment and borrowing (and thus future profits) are influenced not only by current productivity but also by the current shock. A firm with low (high) current profits may choose to invest a lot

(a little) due to good (bad) recent news. This mechanism reduces the correlation between current conditions and firm behavior, allowing the DE model to better match the microdata.

Despite risk neutrality the DE model yields meaningful credit spreads even though the average default frequency is low. This occurs because, in the model, just like in Moody's convention, spreads are measured for firms with moderate (median) credit risk rather than for the mean firm. Given lognormality of the driving productivity process and decreasing returns to scale in the production technology, the mean risk level is lower than the median risk level. Diagnostic expectations by lenders increase the perceived variance of productivity and hence amplify this effect, so the spreads are on average higher with DE than for the RE model.

Overall, this calibration exercise tells us the DE model can better match the micro behavior of firms with respect to investment, leverage, expectations, and spreads. We next assess the ability of this model to account for macroeconomic phenomena, comparing again the DE and RE specifications.

4.3 Diagnosticity and Business Cycle Comovements

Table 7 reports the correlations and volatilities of output growth, leverage growth, investment, mean credit spreads, and default rates computed from an unconditional simulation of the DE model. Table 8 provides the equivalent figures for the RE model, computed from an identical set of macro shocks. Table 9 reports the same aggregate moments in the data. These aggregate business cycle patterns are entirely untargeted in our calibration procedure, with macro TFP shock persistence and volatility set to round values from other papers in the literature to ensure comparability. An examination of the patterns at the macro level still proves useful because it provides insight into the mechanisms at work under diagnosticity. In particular, standard business cycle moments shift considerably with the incorporation of diagnosticity, in a direction that helps to reconcile the model with qualitative patterns in the aggregate data.

Table 7: Business Cycle Moments: Diagnostic (DE) Model

Correlation	Output	Leverage	Investment	Spread
Output	1.0000			
Leverage	0.865	1.0000		
Investment	0.590	0.826	1.0000	
Avg Spread	-0.085	-0.121	-0.153	1.0000
Standard Deviation	0.035	0.016	0.250	0.021

Notes: The table reports the correlations and standard deviations of aggregate output growth, investment growth, leverage growth, and the mean spread. The moments are computed over a 5000-year unconditional simulation in the calibrated DE model with $\theta > 0$.

Table 8: Business Cycle Moments: Rational (RE) Model

Correlation	Output	Leverage	Investment	Spread
Output	1.0000			
Leverage	0.939	1.0000		
Investment	0.749	0.783	1.0000	
Avg Spread	0.121	0.184	0.044	1.0000
Standard Deviation	0.034	0.013	0.151	0.009

Notes: The table reports the correlations and standard deviations of aggregate output growth, investment growth, leverage growth, and the mean spread. The moments are computed over a 5000-year unconditional simulation in the RE model with $\theta = 0$.

Table 9: Business Cycle Moments: Data

Correlation	Output	Leverage	Investment	Spread
Output	1.0000			
Leverage	0.384	1.0000		
Investment	0.844	0.142	1.0000	
Avg Spread	-0.686	-0.310	-0.511	1.0000
Standard Deviation	0.015	0.036	0.074	0.007

Notes: The table reports the empirical correlations and standard deviations of output growth, investment growth, leverage growth, and the mean spread. Output is real GDP from NIPA, investment is private nonresidential fixed investment from NIPA, leverage is the ratio of nonfinancial corporate debt to the net stock of private fixed assets from the Flow of Funds, and the spread is the Moody's BAA-Treasury spread, all at annual frequency from 1986-2017.

The DE model delivers procyclical leverage and credit growth, together with the empirically realistic case of countercyclical spreads. Spreads are countercyclical when the supply of credit expands more than its demand during good times. This does not happen in the RE model because lenders' rational expectations of default are just too stable. As a result, when $\theta = 0$ in the RE model, demand for credit plays a more important role in shaping spreads and creates counterfactual procyclicality.

The DE model generates less comovement of output growth with investment and debt. This is once more due to the fact that in the DE model firms' behavior depends on news, not only on current productivity as in the RE model. This extra source of volatility tends to reduce business cycle comovements in an empirically plausible direction.

Finally, the DE model produces, for an identical macro TFP shock process, more volatility of most business cycle aggregates than the RE model. Given the values for the macro TFP shock process parameters chosen here from the literature, the DE model overpredicts macro volatility relative to the untargeted macro data moments. The lesson here is that belief shifts through diagnosticity represent an amplification device, and a calibration exercise which targeted these

macro moments directly would require smaller exogenous macro TFP shock volatility to generate a given amount of observed business cycle variance.

Before moving further, a discussion of our partial equilibrium or small open economy structure is in order. This structure embeds two maintained assumptions, which we examine in turn. First, we assume a constant interest rate R without general equilibrium stochastic discount factor (SDF) shifts. While the role of the SDF should be explored in future work, we note that the traditional SDF cannot by itself account for our main motivating evidence, namely the predictability of returns on debt and the tendency for crises to arise after credit booms. Furthermore, recent work shows that conventional empirical proxies of the SDF are negatively correlated with survey expectations of returns, raising questions about the validity of the SDF as the driver of expected returns (Greenwood and Shleifer, 2014; Giglio et al., 2019). Assuming a constant interest rate also allows us to avoid making strong assumptions about the nature and cyclicality of SDF movements in practice. Recent work in quantitative macroeconomics studying firm heterogeneity (Khan and Thomas, 2008; Bachmann et al., 2013; Winberry, 2017) emphasizes that general equilibrium SDF movements do not necessarily lead to a dampening of business cycle nonlinearities or investment dynamics if they are structured to produce realistic countercyclicality of real interest rates, while the traditional SDF does in fact dramatically dampen investment dynamics through procyclical real interest rates. Our fixed real interest rate assumption here strikes a middle ground between these alternatives, and is consistent with the evidence that suggests that the real interest rate is only mildly countercyclical (Winberry, 2017).

We also assume a constant wage W. This assumption suggests that the macro volatilities in the DE and RE models above are likely upper bounds. With upward-sloping labor supply and general equilibrium including labor markets, procyclical wage movements would quantitatively dampen the impact of productivity shocks on output, and internalization of such shifts would feed through into dampened investment dynamics. We explore exactly this logic in a general equilibrium extension with labor market clearing in Section 8 below.

The bottom line from our examination of macro moments is that diagnosticity helps account for the countercyclicality of spreads, their volatility, and higher overall variability in the macroeconomy. We next move to assess whether diagnosticity helps account for recurrent and predictable cycles in spreads, investment, and returns on debt.

5 Credit and Investment Cycles under Diagnosticity

We ask two questions of our quantitative model, motivated by the empirical literature on credit cycles and predictable crises. First, can diagnostic expectations help account for financial crises coming after good rather than bad times? Second, can the same mechanism help account for

more regular but still predictable credit and investment cycles? As we show, the answer of both questions is yes, because diagnosticity entails both over-optimism in good times and its systematic subsequent reversal. These features also yield quantitatively reasonable credit cycle movements in the aggregate.

5.1 Financial Crises After Good Times

We define a financial crisis as a period in which the spread grows 1.5 percentage points or more. This threshold corresponds to the 97.5th percentile of spread growth in U.S. data from 1986-2017.[14] We simulate the model for a large number of periods ($T = 500$), identify crises corresponding to this definition, and examine the average dynamics of a variety of aggregate indicators around such events using distributed lag regressions of crisis indicators.

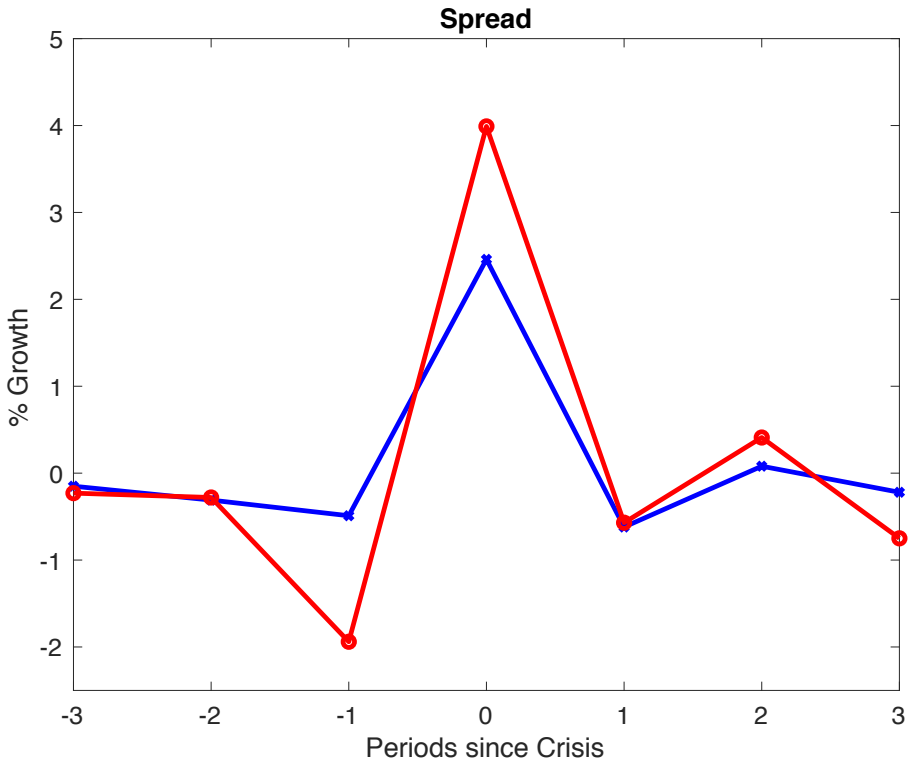


Figure 3: Financial Crisis Dynamics: Spread Growth

Notes: The red line with circles is the DE model, the blue line with × is the RE model. The figure traces out average dynamics around crisis events, defined as periods with greater than 1.5% growth in spreads within an unconditional simulation of the model of 5000 period length.

Figure 3 plots the average path of spread growth around crises in the DE model and in the RE model. By definition, crises are characterized by sharp increases in credit spreads. But there are

[14] Others have used similar ranges in their characterization of financial crises, see Reinhart and Rogoff (2009) and Baron et al. (2019).

two important differences between the two models. First, in the DE model crises occur after good times, i.e., after times with low spreads. This regularity is consistent with the empirical findings of Krishnamurthy and Muir (2016), but does not obtain when $\theta = 0$. Second, the spread dynamics are significantly more pronounced under diagnosticity: during the crisis the spread grows by about six percentage points as opposed to roughly two percentage points in the RE model.

Why do crises occur after periods with low, declining spreads in the model with diagnosticity? The mechanism is illustrated by the TFP dynamics in the top left of Figure 4. In the DE model, the TFP sequence leading to a crisis is "boom and stasis." That is, crises do not come after bad TFP shocks. They come when TFP stops growing after a large boom. Intuitively, during the boom, diagnostic investors are overly optimistic. As a result, they underestimate default risk, charging low spreads. Diagnostic firms are also too optimistic, inducing them to borrow and invest too much. When TFP growth stops, excess optimism wanes, which induces firms to appreciate the risks created by prior over-leveraging. This fragility causes spreads to rise. In the RE model crises are due to sharply different dynamics: they require a prolonged TFP decline. As TFP declines for several periods in a row, the median firm's debt structure becomes less and less sustainable, entailing higher risk of default and higher spreads. This account, though, does not square with the fact that, in reality, crises often come after good times, rather than bad times.

Figure 4 shows how, by creating these boom-bust dynamics, diagnostic expectations affect the behavior of the macroeconomy. Due to the different TFP patterns, output growth, the investment rate, and leverage growth exhibit strong boom-bust dynamics in the DE model, while they exhibit muted dynamics in the RE model. Under diagnosticity, over-leveraging during the boom and de-leveraging during the bust cause large fluctuations in investment. Under rationality, by contrast, firms start to deleverage as the crisis approaches, giving rise to a gradual but small decline in real activity.

To summarize, the DE model creates endogenous financial crises as follows: good news $\Rightarrow$ excess optimism $\Rightarrow$ over-leveraging$\Rightarrow$ systematic reversal in expectations $\Rightarrow$ financial fragility. The ensuing predictability in financial crisis events helps account for the data. We now consider whether the same mechanism can also create the regular and predictability financial cycles documented by López-Salido et al. (2017).

5.2 Predictable Macroeconomic Reversals after Loose Credit Pricing

To assess the predictability of financial and real cycles, we simulate the model for a large number of periods ($T = 500$) and then use the simulated data to run regressions similar to those run by López-Salido et al. (2017)who document predictable macroeconomic reversals after credit spread declines. As a first step, we regress spread growth at $t + 1$ on spread growth at t. This analysis detects any predictable reversal in credit market conditions. As a second step, we regress macro

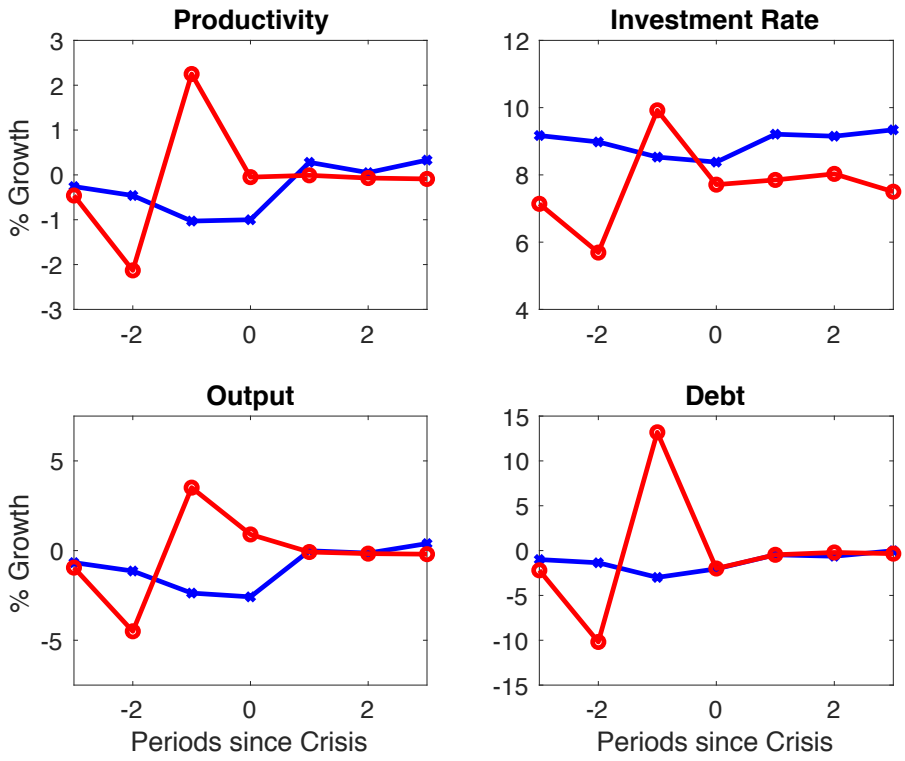


Figure 4: Financial Crisis Dynamics: Macro Aggregates

Notes: In each panel, the red line with circles is the DE model, the blue line with × is the RE model. The top left panel plots macro TFP growth, the top right panel plots the aggregate investment rate, the bottom left plots output growth, and the bottom right plots debt growth. The figure traces out average dynamics around crisis events, defined as periods with greater than 1.5% growth in spreads within an unconditional simulation of the model of 5000 period length.

outcomes of interest at $t+1$ on the contemporaneous spread change that could be predicted using spread growth at t. More precisely, we first regress

$$\Delta \text{Spread}_t = \alpha + \beta \Delta \text{Spread}_{t-1} + \varepsilon_t.$$

We then use the predicted values from this regression $\Delta\widehat{\text{Spread}}_t$ in a second step to evaluate any forecastable patterns in a given macro aggregate X_t by estimating

$$X_t = \delta + \gamma \Delta\widehat{\text{Spread}}_t + \lambda \Delta Y_{t-1} + \eta_t,$$

where following López-Salido et al. (2017) we include lagged output growth ΔY_{t-1}. Table 10 reports $\hat{\beta}$ from the first step (columns (1) and (4)) and $\hat{\gamma}$ from the second step (columns (2)-(3) and (5)-(6)), running such regressions for both the DE and RE models. We consider two outcomes X_t here: aggregate investment growth and realized returns on debt. Greenwood and Hanson (2013) document that realized (excess) bond returns are predictably lower, and sometimes even negative, after periods of low credit spreads. We check whether our model can yield this empirical regularity.

Table 10: Predictable Reversals in the DE Model

Dep. Var.:	Spread Growth$_t$	Investment Growth$_t$	Bond Return$_t$	Spread Growth$_t$	Investment Growth$_t$	Bond Return$_t$
Model:	DE	DE	DE	RE	RE	RE
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted Spread Growth$_t$		-4.777***	-0.070***		10.591***	-0.010
		(1.747)	(0.016)		(3.059)	(0.019)
Spread Growth$_{t-1}$	-0.304***			-0.253***		
	(0.013)			(0.014)		
T	500	500	500	500	500	500

Notes: The table reports a set of regressions run on unconditionally simulated data from the DE model in columns (1)-(3) and the RE model in columns (4)-(6). The underlying macro TFP series is held constant across models. Predicted spread growth is the predicted value from the regression in column (1), i.e., spread growth predicted with lagged spread growth. For the DE model, columns (2)-(3) regress the aggregate investment rate and the average bond return on predicted spread growth. Columns (4)-(6) in the RE model are the analogs of Columns (1)-(3) for the DE model. The investment growth and bond return regressions control for lagged output growth, and all models include constants. Standard errors, in parentheses, are computed from the iid formula. The symbols ***, **, and * imply significance at the 10, 5, and 1% levels.

In column (1) of Table 10, the DE model creates systematic reversals in spread growth, consistent with intuition and with the data. As can be seen from column (4), the RE model also creates some reversals, due to mean reversion in aggregate TFP. However, this effect is smaller than in the DE model because $\theta > 0$ create a reason for reversals even if mean reversion is absent in the RE model: systematic correction of forecast errors.

The difference between the DE and RE models is most pronounced when we consider the predictability of macro outcomes. In column (2), the DE model generates - consistent with the data - a pattern whereby a decline in spreads today is associated with a predictable contraction in future investment. The mechanism is the same as that creating financial crises: excessive optimism leads to over-leveraging, but also to subsequent disappointment and deleveraging. In the RE model, see Column (5) predictability goes the wrong way. Here a reduction in current spreads, due for instance to low TFP and thus stagnating credit demand, is associated with higher future investment, due to mean reversion in productivity.

Diagnosticity is also critical in generating predictability in realized bond returns. Column (3) shows that low current spreads predict low realized bond returns, as in Greenwood and Hanson (2013). This is due to current excess optimism, which leads to neglect of default risk and thus disappointing realized returns. In the RE model, see column (6), realized bond returns are unpredictable. Rational investors will always demand the proper compensation for default risk, which guarantees an average return equal to the constant riskless interest rate.

In sum, by creating excessive optimism in good times, excessive pessimism in bad times, and systematic reversals, the DE model is capable of parsimoniously unifying a range of evidence on financial and real instability.

5.3 Diagnosticity: Demand vs. Supply

Several papers have argued that credit cycles to some extent reflect shocks to either supply or demand for credit. Collin-Dufresn et al. (2001) suggest that the majority of monthly changes in credit spreads are driven by supply or demand factors over and beyond estimates of default risk. Similarly, the predictably low bond returns (Greenwood and Hanson, 2013) and of credit crunches (López-Salido et al., 2017) suggests a role for (predictable) overoptimism by lenders. Our model features demand and supply shocks in the form of predictable deviations from rational expectations, and so it offers a setting in which to quantitatively examine this question.

To do so, we explore how the over-leveraging dynamics that play a central role in the model depend on the structure of the credit market. In particular, we assess to what extent these dynamics arise in a world in which diagnostic firms are disciplined through rational debt pricing.[15]

[15]It is also interesting to examine the case in which firms hold rational expectations but diagnostic credit markets create periods of excessively cheap credit, to which even rational firms may want to adjust. We will explore this

Specifically, we run a simulation of the model in which we set $\theta = 0$ for lenders while keeping θ at the calibrated value for firms. This exercise offers a way to assess which portion of the total variation due to diagnosticity is accounted for by the demand and supply sides. We consider the robustness of the crisis dynamics of Section 4.1, to evaluate the extent to which diagnosticity in credit pricing contributes to the effects we have described.

Figures 5 - 6 below present the crisis plots (still defined as periods in which spread growth is at or above 1.5 percentage points) for the benchmark DE model in which all agents have the same θ and the model in which only firms are diagnostic, while lenders hold rational expectations (they have $\theta = 0$). The message is intuitive: the crisis dynamics lie in between the full diagnostic expectations case and the fully rational expectations case. Intuitively, efficent pricing of credit risk tends to discourage excess leveraging in good times, and both investment and output are smoothed relative to the case of diagnostic lenders (Figure 6). Similarly to the case with rational expectations, crises result from negative TFP shocks.

Importantly, efficient pricing of credit risk does not drive away all excesses. Because overly optimistic diagnostic firms are too eager to borrow, even at the correct interest rate, they put themselves in a fragile position, enhancing the risk of spread reversals, deleveraging, and declines in investment. Now the losses created by diagnostic beliefs fall entirely on firms, not on lenders. This message may be relevant for economic policy: if borrowers' beliefs are diagnostic, a more accurate pricing of risks may not avoid over-borrowing in good times, and while it may reduce total social losses, it may concentrate them on borrowers.

case in the future.

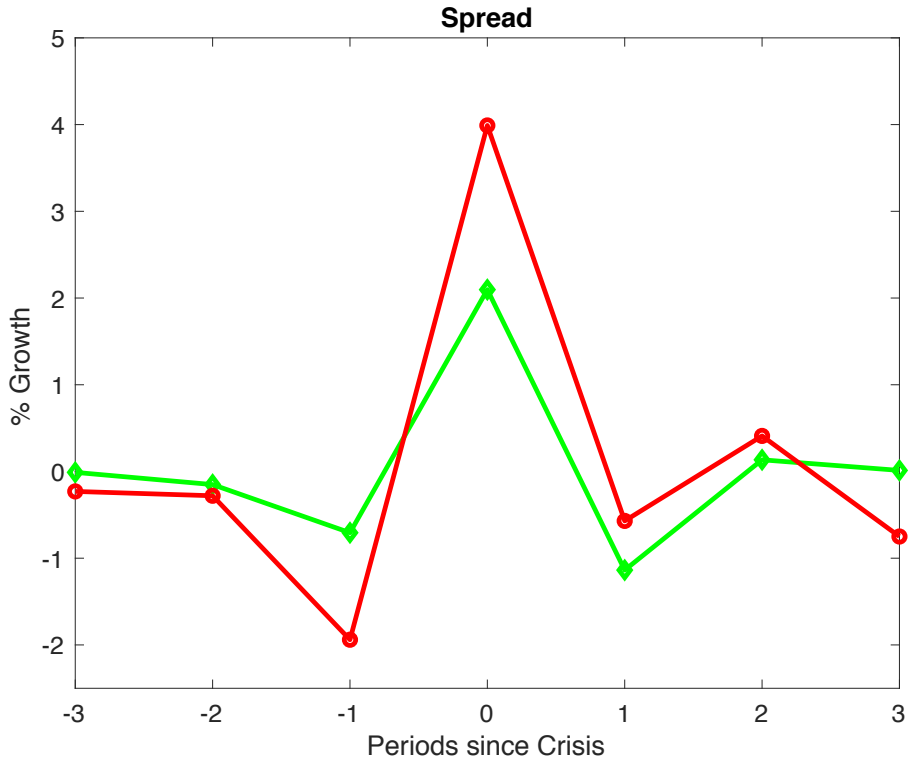


Figure 5: Financial Crisis Dynamics: Spread Growth

Notes: The red line with circles is the baseline DE model, the green line with diamonds is the diagnostic expectations model with rational pricing. The figure traces out average dynamics around crisis events, defined as periods with greater than 1.5% growth in spreads within an unconditional simulation of the model of 5000 period length.

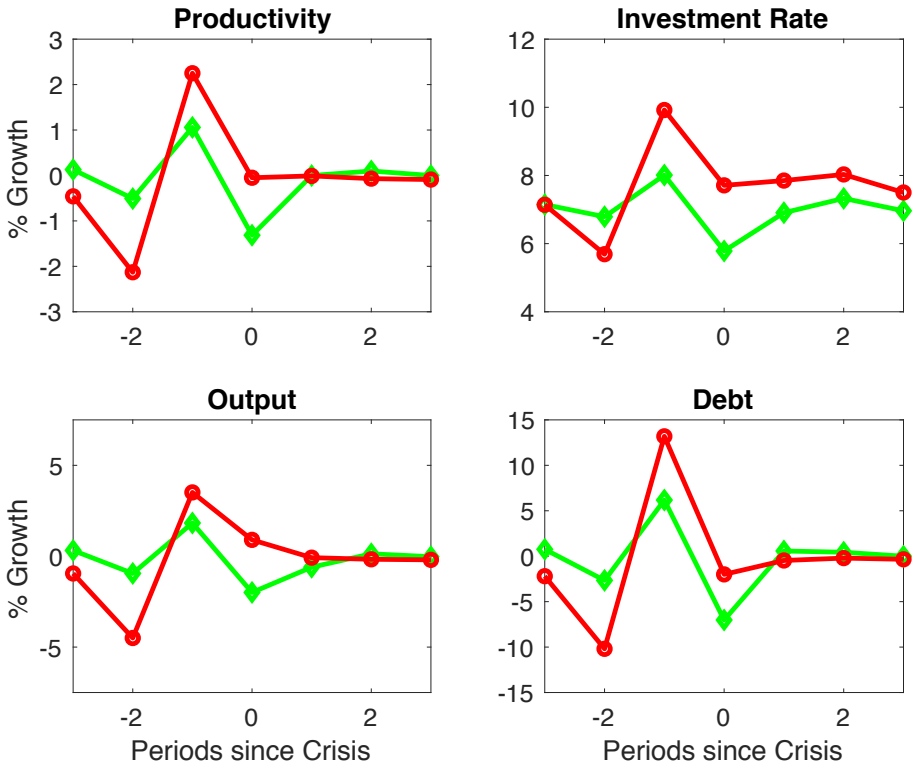


Figure 6: Financial Crisis Dynamics: Macro Aggregates

Notes: In each panel, the red line with circles is the baseline DE model, the green line with diamonds is the diagnostic expectations model with rational pricing. The top left panel plots macro TFP growth, the top right panel plots the aggregate investment rate, the bottom left plots output growth, and the bottom right plots debt growth. The figure traces out average dynamics around crisis events, defined as periods with greater than 1.5% growth in spreads within an unconditional simulation of the model of 5000 period length.

6 Diagnosticity in the 2008 US Crisis

The US financial crisis of 2008 shares several features of the financial instability analysed in the previous section. First, it materialized after a period of sustained economic growth and of unusually low credit spreads. Second, it witnessed a leverage and an investment cycle, with strong expansions in the pre crisis periods followed by sharp contractions. A growing body of work, both theoretical and empirical, places extrapolative beliefs center stage in explaining this episode (Gennaioli and Shleifer, 2018). This section seeks to assess the extent to which diagnosticity can shed light on the macroeconomic causes and consequences of this episode.

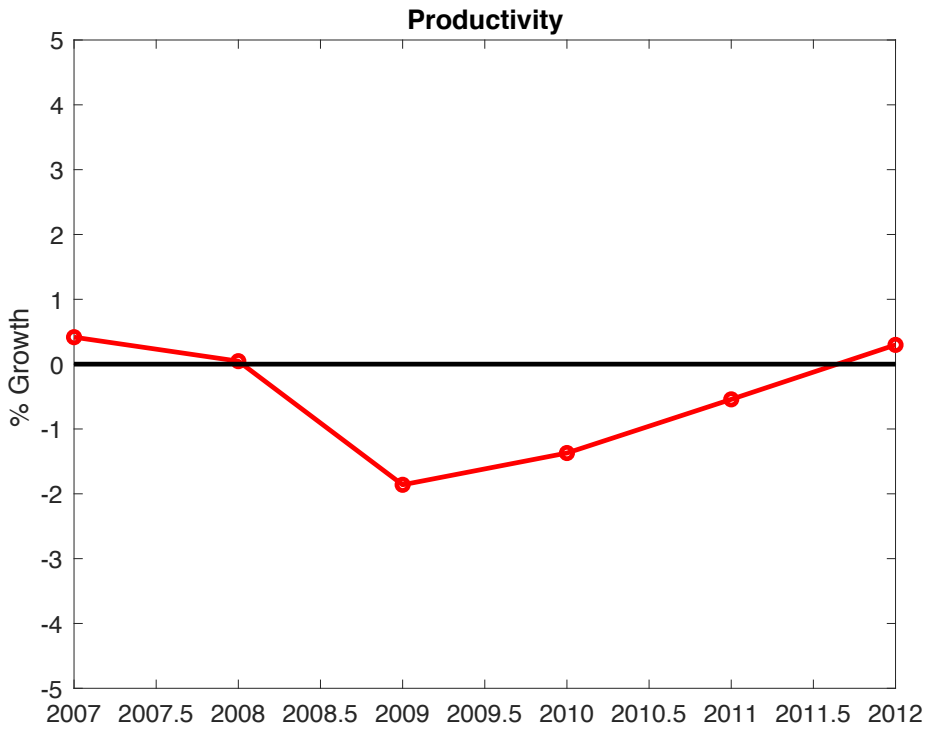


Figure 7: Historical Decomposition: Matching the 2008 Crisis

Notes: The figure plots the dynamics of macro TFP required in the DE model in order to exactly match empirical macro investment growth over the period surrounding the 2008 US financial crisis, as described in the main text. This process for macro TFP is computed according to the historical decomposition procedure described in Appendix A.

The logic of our exercise goes as follows. We use our nonlinear calibrated DE model to determine an aggregate TFP sequence that exactly matches the dynamics of US aggregate investment growth from 2007 to 2012. This is a computationally intensive procedure detailed in Appendix A, and the resulting series is plotted in Figure 7. We then assess the implications of this aggregate TFP series for output, investment, debt, and credit spreads, comparing them in each case to empirical quantities in Figure 8. To tease out the effects of diagnosticity, we also assess the

implications of the *same* aggregate TFP sequence for the RE model.

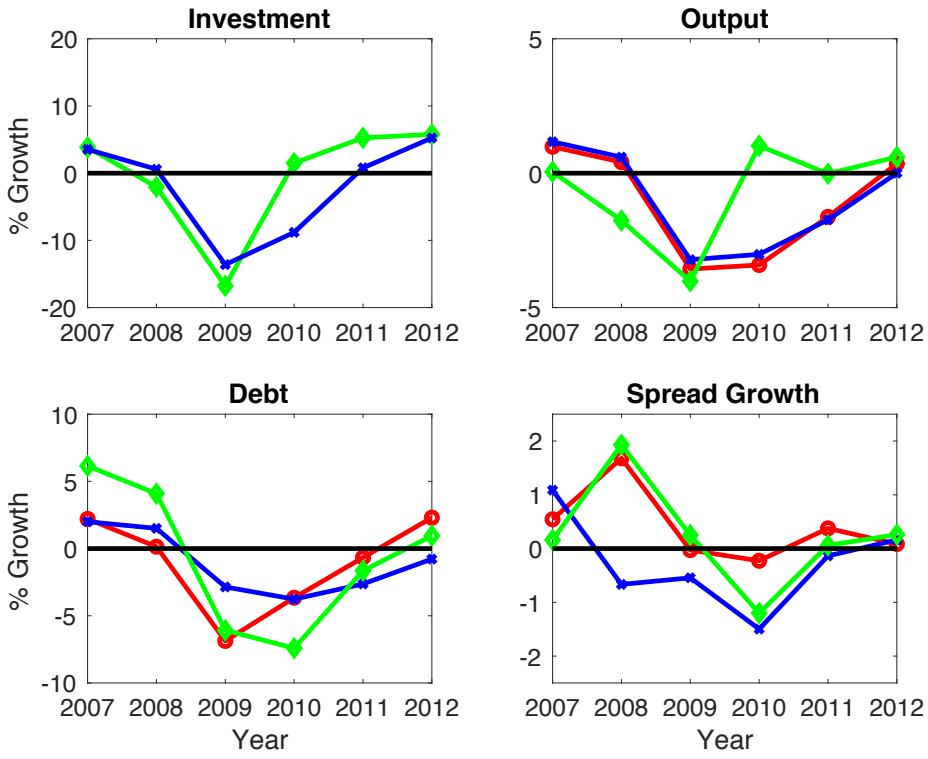


Figure 8: Historical Decomposition: Macro Aggregates

Notes: Throughout the figure, red lines with circles denote the DE model. The blue lines with × denote the RE model. The green lines with diamonds denote the data. In all cases, the theoretical series are a function of the same aggregate macro TFP series in Figure 7 chosen to exactly match the empirical investment dynamics (the green line in the top left panel) when fed through the DE model.

We solve the model at a fairly coarse annual frequency and lack sectoral variation such as a distinction between real estate and other sectors. While financial trouble began to manifest itself in housing-related markets in 2007, broader macro investment did not begin to fall appreciably until later. So our DE model, with shocks chosen to match investment dynamics, requires negative growth in aggregate TFP roughly equal to -1.5% in 2009 then gradually returning to zero in the following years, plotted in Figure 7. This is a significant but not dramatic productivity decline. At this TFP path the DE model is by construction capable of matching the actual investment growth dynamics, in the top left panel of Figure 8. Yet, a comparison with the implied investment growth path in the RE model already reveals the role of diagnosticity. As evident in the top left of Figure 8 the RE model somewhat underpredicts the extent of the drop in investment and dramatically underpredicts the recovery of investment growth afterwards. The reason for this failure is intuitive: diagnostic firms and investors over-react to the initial shock, which amplifies the adverse impact on investment and deleveraging. Later on, however, this excess pessimism wanes, so diagnostic

firms are significantly less leveraged than in the RE model. As a result, recovery is faster.

Both the DE and RE models feature output growth which declines by similar magnitudes as the data, although the larger drop in investment in the DE model naturally implies a lower path for output growth during the later stages of the recovery. The fact that output differences between DE and RE is small is expected, given that in both models the business cycle is driven by the same TFP series. One interesting area for future work is to add demand side frictions or capital utilization mechanisms in the DE model.

Diagnosticity also helps explain the behavior of debt and spreads. The bottom left panel of Figure 8 shows the dynamics of debt growth. Debt dynamics in the DE model track reality, displaying leverage growth before the crisis, massive collapse in 2008, and swift recovery afterwards. The RE model cannot account for these reversals, featuring a much smoother and slower decline and recovery than reality. Similarly, diagnosticity can account for the swift hike of credit spreads in 2009 and for their subsequent fast decline, while the RE model cannot: here the spread actually declines during the crisis, due to reduced demand for capital by firms.

This exercise shows the potential of diagnostic expectations for accounting for some of the salient features of the crisis events, but also the limitations of our current setting. First of all, we abstract from the housing bubble. As a result, to capture the events in the crisis we need a negative TFP shock. A model accounting for the role of housing both as collateral in financial transactions and as a means of financing household consumption might, perhaps more realistically, create a crisis from the gradual deflation of the housing bubble itself, without requiring any negative TFP shocks. Second, we abstract from intermediary leverage. The importance in reality of the financial constraints of intermediaries – the lenders in our model – is also likely to severely limit the explanatory power of our model, in which these constraints are absent. When lenders' financial fragility is taken into account, even a mild reversal in optimism may create a strong disruption in the supply of capital for all firms, include the more productive ones, further reducing the drop in TFP required to account for the crisis.

Diagnostic expectations can be combined with these richer mechanisms, and in fact Maxted (2019) examines a model along these lines, incorporating diagnostic beliefs in a model of intermediary leverage. Developing more realistic models combining financial frictions and housing with diagnostic beliefs is an important avenue for future work.

7 Heightened Sensitivity of the Economy during Booms

Recent work on investment dynamics over the business cycle (Bachmann et al., 2013; Winberry, 2017; Bloom et al., 2018) suggests that the business cycle - and investment in particular - exhibits more sensitivity to shocks during booms than during normal times. This type of state-dependence

or non-linearity arises naturally in the DE model. The mechanism proves simple and related to the financial crisis dynamics detailed above. Sharp overinvestment and high leveraging by firms during booms generate sensitivity to even moderately small negative shocks. In contrast, the RE model lacks the overinvestment or leveraging required to generate heightened fragility during booms.

To illustrate the state-dependence of the DE model, we report impulse responses of the DE and RE models to negative macro TFP shocks occurring after two alternative preceding histories: neutral news and good news. Since the model is nonlinear, we implement these impulse responses using the Generalized Impulse Response concept outlined by Koop et al. (1996) and detailed in Appendix A, simulating a large number of pairs of "shocked" vs "unshocked" economies, and taking the average difference across the two as our impulse response series.

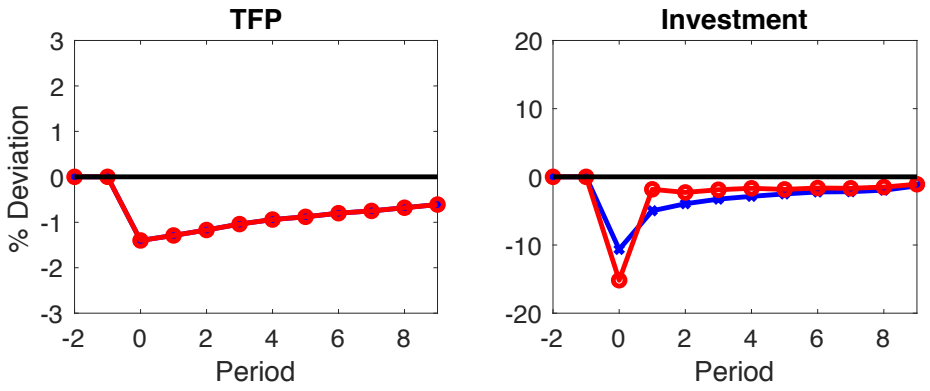


Figure 9: Impulse Response to a Negative TFP Shock: Normal Times

Notes: The figure plots productivity and investment impulse responses in the DE model (red lines with circles) and RE model (blue line with ×) to a negative shock to productivity. The impulse responses are nonlinear generalized impulse responses computed according to the procedure laid out in Appendix A and yield a shock size equal to one-standard deviation of the macro TFP process.

Figure 9 plots the path of productivity from a single negative macro TFP shock after neutral times in the DE vs RE models, identical across the two environments and scaled to equal one standard deviation of the macro TFP shock process in magnitude. After this single negative shock occurring during normal times, investment declines in both models. Although the investment drop is slightly sharper in the DE model with overreaction, the difference in magnitudes or shapes of the impulse response across the DE vs RE models is not severe.

In contrast, Figure 10 plots the impact of a negative TFP shock in both models coming directly after a positive TFP shock. In this case, because diagnostic firms overinvest after the preceding positive shock, the reversal in investment in the DE model is large and severe. Instead, the response of investment to the negative shock in the RE model exhibits little difference from the negative shock in neutral times shown in Figure 9. In other words, the DE model's investment

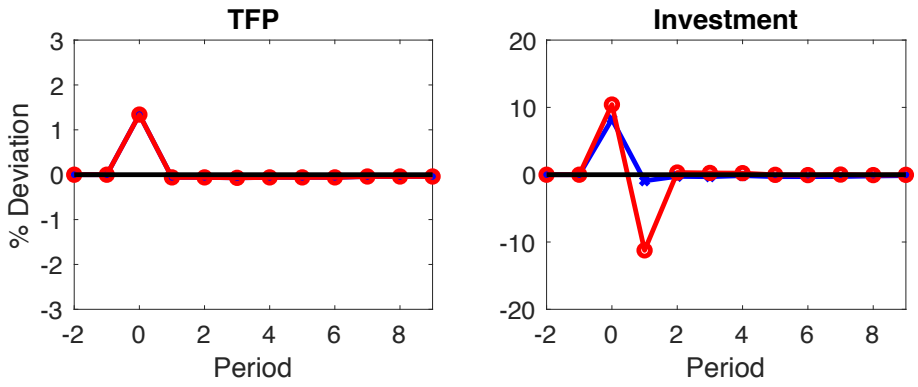


Figure 10: Impulse Response to a Negative TFP Shock: After Good News

Notes: The figure plots productivity and investment impulse responses in the DE model (red line with circles) and RE model (blue line with $\times$) to a single negative shock to productivity following a single positive shock to productivity. The impulse responses are nonlinear generalized impulse responses computed according to the procedure laid out in Appendix A and yield a shock size in each direction equal to one-standard deviation of the macro TFP process.

behavior is clearly more sensitive to shocks after good times than bad, but the RE model shows little state-dependence.

To summarize, diagnosticity generates state-dependence or nonlinearity in the macroeconomy. During booms, overoptimism on the part of firms and lenders generates more investment sensitivity than during normal times. The DE model generates these patterns without the incorporation of mechanisms such as lumpy capital adjustment costs or uncertainty shocks sometimes employed to generate such business cycle nonlinearities. Instead, diagnostic expectations by themselves prove to be a distinct and powerful source of state-dependence.

8 General Equilibrium

In this section we endogenize the wage W, extending our analysis from the partial equilibrium economy considered so far towards general equilibrium. As discussed above, empirically realistic countercyclical shifts in real interest rates in this class of heterogeneous firms models are likely to amplify rather than dampen investment dynamics. By contrast, anticipated procyclical wage shifts dampen movements in anticipated marginal products of capital and hence push against volatility or nonlinearity in investment. We therefore conservatively consider a representative household which is risk neutral in consumption C, fixing the real interest rate R, but allowing for a market-clearing wage W. Period utility is given by

$$U(C,N) = C - \frac{\omega}{1+\frac{1}{\lambda}} N^{1+\frac{1}{\lambda}},$$

where the disutility of labor is governed by $\omega > 0$ and the elasticity of labor supply is given by $\lambda > 0$. In this environment, real interest rates are pinned down by the inverse of the household's subjective discount rate $0 < \beta < 1$, i.e., $R = \frac{1}{\beta}$.

Let $\mu(s, k, b)$ be the cross-sectional distribution of exogenous states s, capital k, and debt b. The aggregate state is (μ, A, ε_A). Market clearing in the labor market implicitly defines wages $W(\mu, A, \varepsilon_A)$ through

$$\left(\frac{W}{\omega}\right)^{\lambda} = \int n(s, k, b|W) d\mu(s, k, b),$$

where the left hand side is the household's closed-form labor supply and the right hand side reflects labor demand generated by the current cross-sectional distribution of firm states μ.

As usual in this class of heterogeneous firms models with anticipated macro shocks, general equilibrium creates two computational challenges (Terry, 2017a). First, the macro state (μ, A, ε_A) is intractable because μ is a distribution. Second, the mapping $W(\mu, A, \varepsilon_A)$ is a complicated implicit object which must be consistent with the firm-level decisions embedded in market clearing.

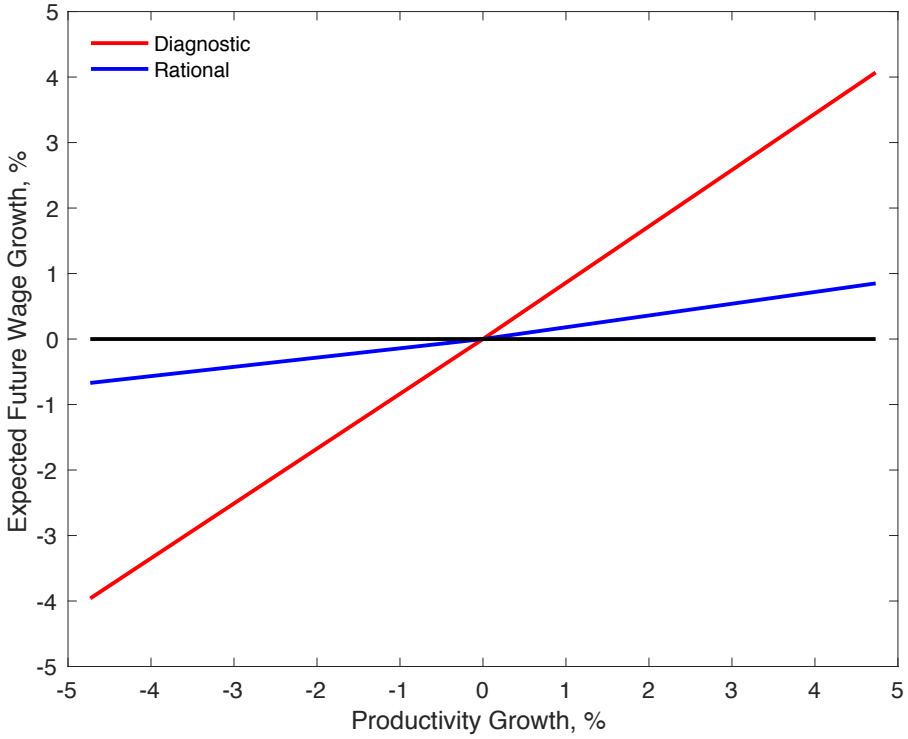


Figure 11: Extrapolated Wage Growth in General Equilibrium

Notes: The figure plots expected future wage growth $\mathbb{E}[\Delta \log W' | \Delta \log A]$ as a function of today's macro TFP growth in the DE model (red) and RE model (blue) under general equilibrium.

We follow a computational approach tailored to our problem and detailed in Appendix A. To briefly summarize, our approach is to replace the aggregate state $(\mu_t, A_t, \varepsilon_{At})$ with a history of

aggregate shocks $(A_t, A_{t-1}, ..., A_{t-K})$ up to some truncated lag length, nonparametrically storing predictions of the wage W given each shock history. We then follow an outer loop/inner loop approach, guessing a wage mapping, solving and simulating the model, and updating the wage predictions until convergence. Our solution technique proves tractable and accurate in practice, as also detailed in Appendix A. We parameterize the model based on the calibrated parameters from Table 5. We further assume a conservative Frisch elasticity of labor supply of $\lambda = 0.5$ and choose β to deliver the same fixed 4% annual real interest rate as considered above.

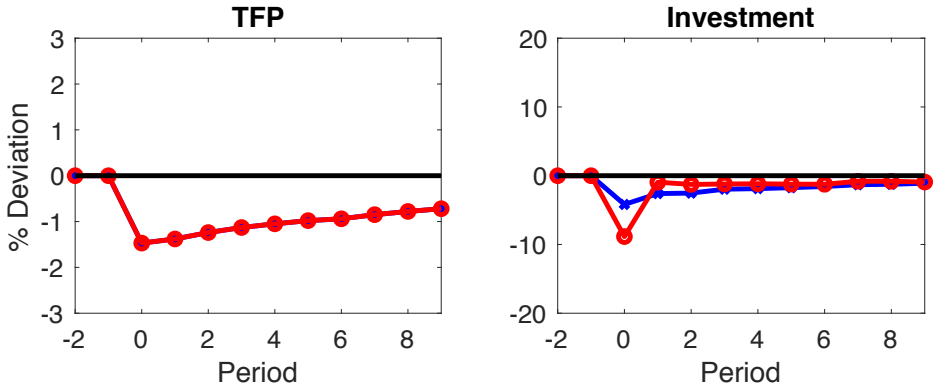


Figure 12: GE IRF to a Negative TFP Shock: Normal Times

Notes: The figure plots GE productivity and investment impulse responses in the DE model (red lines with circles) and RE model (blue line with ×) to a negative shock to productivity. The impulse responses are nonlinear generalized impulse responses computed according to the procedure laid out in Appendix A and yield a shock size equal to one-standard deviation of the macro TFP process.

To gauge the impact of labor market clearing on firms' intertemporal choices, Figure 11 presents the predicted future growth in wages as a function of current macro TFP growth. To the extent that firms expect higher offsetting wage growth after growth in macro TFP, their demand for and hence investment in complementary capital will grow less dramatically. Figure 11 documents that in both the DE and RE models, expected future wage growth perceived by firms slopes upward, dampening investment dynamics. However, due to diagnostic beliefs, firms in the DE model expect even higher wage growth than RE firms after positive growth in TFP. Apparently, the discipline on investment imposed by endogenous wages proves even stronger for the diagnostic than for the rational expectations model.

We now investigate whether the state-dependence or nonlinearity generated by diagnostic beliefs in this model survives general equilibrium. Figure 12 plots the response of the general equilibrium economy to a negative macro TFP shock during normal times, with qualitatively similar drops in investment for both the DE and RE models. In contrast, Figure 13 plots the impact of a negative macro TFP shocks during good times, i.e., after a preceding positive macro TFP shock. Because of the overoptimism, overleveraging, and overinvestment of firms after good news, the

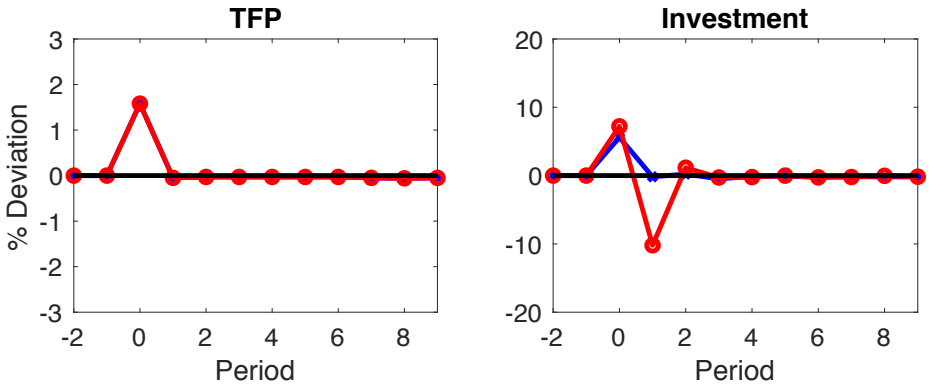


Figure 13: GE IRF to a Negative TFP Shock: After Good News

Notes: The figure plots GE productivity and investment impulse responses in the DE model (red lines with circles) and RE model (blue line with ×) to a negative shock to productivity following a single positive shock. The impulse responses are nonlinear generalized impulse responses computed according to the procedure laid out in Appendix A and yield a shock size equal to one-standard deviation of the macro TFP process.

DE economy proves more fragile and responsive to negative shocks than the RE economy, with a plunge in investment exceeding that seen in normal times. Comparison of these responses to their partial equilibrium counterparts in Figures 9-10 above reveals a modest quantitative dampening of investment, as expected. Overall, the role of diagnostic beliefs in driving macro-financial dynamics survives general equilibrium price shifts.

9 Conclusion

Macro fragility naturally arises in a canonical business cycle model as a result of micro-founded deviations from rational expectations by individual firms and creditors. Business cycle dynamics in our quantitative neoclassical business cycle model incorporating realistic micro-level expectations and heterogeneity prove more volatile, less stable, and feature sharp crises with rapidly worsening credit conditions, deleveraging, and sharp recessions. Such crises occur after good times with expansion of credit and low spreads. Since a rational expectations model fails to capture such realistic credit cycle dynamics, realistic modeling of expectations may provide a useful tool for understanding macro-financial fluctuations.

References

Alfaro, Ivan, Nicholas Bloom, and Xiaoji Lin (2018), "The finance uncertainty multiplier." Technical report, National Bureau of Economic Research.

Arellano, Cristina, Yan Bai, and Patrick J. Kehoe (2019), "Financial frictions and fluctuations in volatility." *Journal of Political Economy*, 127, 2049–2103.

Bachmann, Rüdiger, Ricardo J Caballero, and Eduardo MRA Engel (2013), "Aggregate implications of lumpy investment: new evidence and a dsge model." *American Economic Journal: Macroeconomics*, 5, 29–67.

Baron, Matthew, Emil Verner, and Wei Xiong (2019), "Salient crises, quiet crises." Working paper.

Baron, Matthew and Wei Xiong (2017), "Credit expansion and neglected crash risk." *The Quarterly Journal of Economics*, 132, 713–764.

Barrero, Jose M. (2018), "The micro and macro of manager beliefs." Working Paper.

Bernanke, Ben and Mark Gertler (1989), "Agency costs, net worth, and business fluctuations." *The American Economic Review*, 79, 14–31.

Bernanke, Ben S, Mark Gertler, and Simon Gilchrist (1999), "The financial accelerator in a quantitative business cycle framework." *Handbook of macroeconomics*, 1, 1341–1393.

Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho (2019), "Survey data and subjective beliefs in business cycle models." Working paper.

Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma (2020), "Belief distortions and macroeconomic fluctuations." NBER working paper 27406.

Bloom, Nicholas (2009), "The impact of uncertainty shocks." *Econometrica*, 77, 623–685.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry (2018), "Really uncertain business cycles." *Econometrica*, 86, 1031–1065.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer (Forthcominga), "Memory and representativeness." *Psychological Review*. Working paper.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2016), "Stereotypes." *The Quarterly Journal of Economics*, 131, 1753–1794.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2019a), "Beliefs about gender." *American Economic Review*, 109, 739–73.

Bordalo, Pedro, Nicola Gennaioli, Rafael LaPorta, and Andrei Shleifer (2019b), "Diagnostic expectations and stock returns." *Journal of Finance*, 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (Forthcomingb), "Over-reaction in macroeconomic expectations." *American Economic Review*. NBER working paper.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018), "Diagnostic expectations and credit cycles." *The Journal of Finance*, 73, 199–227.

Borio, Claudio E.V. and Philip W Lowe (2002a), "Assessing the risk of banking crises." *BIS Quarterly Review*, 43–54. BIS Quarterly Review.

Borio, Claudio E.V. and Philip W Lowe (2002b), "Asset prices, financial and monetary stability: Exploring the nexus." BIS Working Paper No 114.

Bouchaud, Jean-philippe, Philipp Krueger, Augustin Landier, and David Thesmar (2019), "Sticky expectations and the profitability anomaly." *The Journal of Finance*, 74, 639–674.

Brunnermeier, Markus K, Gary Gorton, and Arvind Krishnamurthy (2012), "Risk topography." *NBER Macroeconomics Annual*, 26, 149–176.

Brunnermeier, Markus K and Yuliy Sannikov (2014), "A macroeconomic model with a financial sector." *American Economic Review*, 104, 379–421.

Casscells, Ward, Arno Schoenberger, and Thomas B Graboys (1978), "Interpretation by physicians of clinical laboratory results." *New England Journal of Medicine*, 299, 999–1001.

CBO (2017), "International comparisons of corporate income tax rates." Exhibit 7.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno (2014), "Risk shocks." *American Economic Review*, 104, 27–65.

Collin-Dufresn, Pierre, Robert S Goldstein, and J Spencer Martin (2001), "The determinants of credit spread changes." *The Journal of Finance*, 56, 2177–2207.

Dávila, Eduardo and Anton Korinek (2017), "Pecuniary externalities in economies with financial frictions." *The Review of Economic Studies*, 85, 352–395.

Den Haan, Wouter J (2010), "Comparison of solutions to the incomplete markets model with aggregate uncertainty." *Journal of Economic Dynamics and Control*, 34, 4–27.

Eggertsson, Gauti B and Paul Krugman (2012), "Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach." *The Quarterly Journal of Economics*, 127, 1469–1513.

Farhi, Emmanuel and Iván Werning (2016), "A theory of macroprudential policies in the presence of nominal rigidities." *Econometrica*, 84, 1645–1704.

Fostel, Ana and John Geanakoplos (2014), "Endogenous collateral constraints and the leverage cycle." *Annual Review of Economics*, 6, 771–799.

Fuster, Andreas, David Laibson, and Brock Mendel (2010), "Natural expectations and macroeconomic fluctuations." *Journal of Economic Perspectives*, 24, 67–84.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer (2016), "Expectations and investment." *NBER Macroeconomics Annual*, 30, 379–431.

Gennaioli, Nicola and Andrei Shleifer (2010), "What comes to mind." *The Quarterly journal of economics*, 125, 1399–1433.

Gennaioli, Nicola and Andrei Shleifer (2018), *A crisis of beliefs: Investor psychology and financial fragility*. Princeton University Press.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (2019), "Five facts about beliefs and portfolios." Working paper.

Gilchrist, Simon and Egon Zakrajšek (2012), "Credit spreads and business cycle fluctuations." *American Economic Review*, 102, 1692–1720.

Gilchrist, Simon, Egon Zakrajšek, and Jae Sim (2014), "Uncertainty, financial frictions, and investment dynamics." Working paper.

Gourio, Francois (2013), "Credit risk and disaster risk." *American Economic Journal: Macroeconomics*, 5, 1–34.

Gourio, Francois and Leena Rudanko (2014), "Customer Capital." *Review of Economic Studies*.

Greenwood, Robin and Samuel G Hanson (2013), "Issuer quality and corporate bond returns." *The Review of Financial Studies*, 26, 1483–1525.

Greenwood, Robin, Samuel G. Hanson, and Lawrence J. Jin (2019), "Reflexivity in credit markets." Working paper.

Greenwood, Robin and Andrei Shleifer (2014), "Expectations of returns and expected returns." *The Review of Financial Studies*, 27, 714–746.

Gu, Chao, Fabrizio Mattesini, Cyril Monnet, and Randall Wright (2013), "Endogenous credit cycles." *Journal of Political Economy*, 121, 940–965.

Guerrieri, Veronica and Guido Lorenzoni (2017), "Credit crises, precautionary savings, and the liquidity trap." *The Quarterly Journal of Economics*, 132, 1427–1467.

Gulen, Huseyin, Mihai Ion, and Stefano Rossi (2019), "Credit cycles and corporate investment." Working paper.

Hamilton, David T. and Richard Cantor (2005), "Rating transitions and defaults conditional on rating outlooks revisited: 1995-2005." Technical report, Moody's Investor Service.

Hansen, Lars Peter and Thomas J Sargent (2001), "Acknowledging misspecification in macroeconomic theory." *Review of Economic Dynamics*, 4, 519–535.

He, Jie Jack and Xuan Tian (2013), "The dark side of analyst coverage: The case of innovation." *Journal of Financial Economics*, 109, 856–878.

He, Zhiguo and Arvind Krishnamurthy (2019), "A macroeconomic framework for quantifying systemic risk." *American Economic Journal: Macroeconomics*, 11, 1–37.

Hennessy, Christopher A and Toni M Whited (2007), "How Costly is External Financing? Evidence from a Structural Estimation." *Journal of Finance*, 62, 1705–1745.

Jaimovich, Nir and Sergio Rebelo (2007), "Behavioral theories of the business cycle." *Journal of the European Economic Association*, 5, 361–68.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2013), "When credit bites back." *Journal of Money, Credit and Banking*, 45, 3–28.

Kahana, Michael Jacob (2012), *Foundations of human memory*. OUP USA.

Kahneman, Daniel and Amos Tversky (1972), "Subjective probability: A judgment of representativeness." *Cognitive psychology*, 3, 430–454.

Khan, Aubhik and Julia K Thomas (2008), "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics." *Econometrica*, 76, 395–436.

Khan, Aubhik and Julia K Thomas (2013), "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Journal of Political Economy*, 121, 1055–1107.

Kiyotaki, Nobuhiro and John Moore (1997), "Credit cycles." *Journal of political economy*, 105, 211–248.

Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996), "Impulse response analysis in nonlinear multivariate models." *Journal of Econometrics*, 74, 119–147.

Korinek, Anton and Alp Simsek (2016), "Liquidity trap and excessive leverage." *American Economic Review*, 106, 699–738.

Krishnamurthy, Arvind and Tyler Muir (2016), "How credit cycles across a financial crisis." Working paper.

López-Salido, David, Jeremy C Stein, and Egon Zakrajšek (2017), "Credit-market sentiment and the business cycle." *The Quarterly Journal of Economics*, 132, 1373–1426.

Lorenzoni, Guido (2008), "Inefficient credit booms." *The Review of Economic Studies*, 75, 809–833.

Lucas, Robert E (1976), "Econometric policy evaluation: A critique." In *Carnegie-Rochester conference series on public policy*, volume 1, 19–46.

Maxted, Peter (2019), "Systemic risk in a macroeconomic model with diagnostic expectations." Working paper.

Mian, Atif, Amir Sufi, and Emil Verner (2017), "Household debt and business cycles worldwide." *The Quarterly Journal of Economics*, 132, 1755–1817.

Minsky, Hyman P (1977), "The financial instability hypothesis: An interpretation of keynes and an alternative to standard theory." *Challenge*, 20, 20–27.

Reinhart, Carmen M. and Kenneth S. Rogoff (2009), *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Rognlie, Matthew, Andrei Shleifer, and Alp Simsek (2018), "Investment hangover and the great recession." *American Economic Journal: Macroeconomics*, 10, 113–53.

Saporta-Eksten, Itay and Stephen J. Terry (2018), "Short-term shocks and long-term investment." Working paper.

Schularick, Moritz and Alan M Taylor (2012), "Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008." *American Economic Review*, 102, 1029–61.

Shleifer, Andrei and Robert W Vishny (1992), "Liquidation values and debt capacity: A market equilibrium approach." *The Journal of Finance*, 47, 1343–1366.

Simsek, Alp (2013), "Belief disagreements and collateral constraints." *Econometrica*, 81, 1–53.

Stein, Jeremy C (2012), "Monetary policy as financial stability regulation." *The Quarterly Journal of Economics*, 127, 57–95.

Strebulaev, Ilya A and Toni M Whited (2012), "Dynamic Models and Structural Estimation in Corporate Finance." *Foundations and Trends in Finance*, 6, 1–163.

Tauchen, George (1986), "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics letters*, 20, 177–181.

Terry, Stephen J (2017a), "Alternative methods for solving heterogeneous firm models." *Journal of Money, Credit and Banking*, 49, 1081–1111.

Terry, Stephen J. (2017b), "The macro impact of short-termism." Working paper.

Tversky, Amos and Daniel Kahneman (1983), "Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment." *Psychological review*, 90, 293.

Winberry, Thomas (2017), "Lumpy investment, business cycles, and stimulus policy."

Young, Eric R (2010), "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell–Smith Algorithm and Non-Stochastic Simulations." *Journal of Economic Dynamics and Control*, 34, 36–41.

Appendices

A Model

A.1 Solving the Model

The computational algorithm involves iteration on an outer loop (related to debt pricing) and an inner loop (related to firm policies). Before solving the model, we discretize the state space $(s, k, b) = (z, \eta_z, A, \eta_A, k, b)$ into $n_z \times n_z \times n_A \times n_A \times n_k \times n_b$ grid points, with log-linear spacing. We then discretize the rational and perceived diagnostic transitions of the exogenous states according to Tauchen (1986). The computational algorithm - following Strebulaev and Whited (2012) - proceeds as follows:

Start outer loop.

1. Guess a default policy $df^{\theta}(s, k, b)$, and compute the implied debt prices $q^{\theta}(s, k, b)$ according to the lenders diagnostic zero-profit condition.

   Start inner loop.

   (a) Given the debt prices $q^{\theta}(s, k, b)$ and default policy $df^{\theta}(s, k, b)$, solve the diagnostic firm's Bellman equations $V^{\theta}(s, k, b)$, $V^{\theta}_{ND}(s, k, b)$ for the implied optimal policies for investment and debt issuance $k'^{\theta}(s, k, b)$, $b'^{\theta}(s, k, b)$. Use standard discrete-state, discrete-policy dynamic programming policy iteration to do so.

2. Compute updated default policies $df^{\theta}(s, k, b)$ according to the limited liability default condition defining V^{θ}.

3. If the updated default states are identical to the initially guessed default states, exit. If not, then go to top and restart.

We implement this computationally intensive algorithm in heavily parallelized Fortran. Table 1 reports the value of several dimensions used for the baseline solution of the model.

A.2 Simulating the Model

After the model is solved, we unconditionally simulate the model by drawing exogenous uniform random shocks and combining this information with the transition matrix for macro TFP to simulate the aggregate process for A_t for some periods $t = 1, ..., T^{sim} + T^{erg}$. At the micro level,

Table 1: Computational Choices

Quantity	Description	Value
T^{sim}	Simulated periods	5000
T^{erg}	Initially discarded periods	250
N^{firm}	Number of firms	2500
N^{IRF}	Number of IRF economies	1000
T^{IRF}	Length of IRF economies	75
T^{IRF}	Length of historical decomposition	12
n_z	Micro productivity grid size	5
n_A	Macro productivity grid size	5
n_k	Capital grid size	35
n_b	Debt grid size	35

Notes: The table reports various computational values used in discretizing and solving the model.

we simulate the model "non-stochastically" according to the method of Young (2010), i.e., we store the dynamics of the weight of the cross-sectional distribution at each discretized point in the state space (s, k, b) rather than simulating a large number of firms. Note that when simulating the model, all aggregate shocks and distributional dynamics are determined according to the *rational* or true representations of the driving process, even though debt pricing and firm polices may involve diagnostic expectations.

With the simulated distribution in hand for each period, aggregate series of interest are simply weighted sums of micro-level outcomes across this distribution, discarding the first T^{erg} periods to remove the influence of initial conditions. Note that we do in fact simulate a number of individual firms N^{firm} for the purpose of computing moments, but this is not a step required for the purpose of solving the model or simulating within-period business cycle aggregates.

With the solution algorithm above in hand, we calibrate the model by changing only the targeted parameters in each moment calculation iteration, keeping the aggregate shocks unchanged. We minimize the sum of squared percentage deviation of simulated vs data targeted moments, by employing a global stochastic optimization routine.

A.3 Computing Impulse Responses

Our approach to impulse response calculation in this nonlinear context follows Koop et al. (1996), i.e., we compute nonlinear generalized impulse responses. To understand the impact of a given sequence of shocks, we perform the following:

1. For a large number N^{irf} of economies of length T^{irf}, simulate two different versions of the simulation, the "shock" and "no shock" versions. For each economy and each version, we simulate the macro TFP process by first drawing T^{irf} uniform shocks for comparison with the macro TFP transition matrix. Then, simulate both versions unconditionally using identical macro TFP shocks until period $T^{shock} < T^{irf}$.

2. From period T^{shock} and continuing as long as the desired sequence of exogenous innovations you wish to impose lasts, impose a number of periods of certain pre-determined innovations in productivity for the "shock" case, while continuing to simulate the "no shock" economy unconditionally.

3. After the imposed shocks sequence is complete, simulate macro TFP in both economies as normal.

4. After the macro TFP process is determined for each pair of economies, compute the business cycles aggregates of interest in each economy, period, and version by using the simulation approach outlined above.

5. If business cycle aggregate $X_{i,t}^{shock}$ is series X in economy i in period t in the shock case, and $X_{i,t}^{noshock}$ is series X in economy i in period t in the no shock case, then define the impulse response to the predetermined sequence of innovations as

$$IRF_t^X = \frac{1}{N^{irf}} \sum_{i=1}^{N^{irf}} \frac{X_{i,t}^{shock} - X_{i,t}^{noshock}}{X_{i,t}^{noshock}}.$$

The main text's set of impulse response figures reports the series IRF^X for the indicated macro-financial aggregates. Note, however, that the impulse responses presented in the text are scaled to equal an exact shock size, while the productivity grid in the model varies discretely. We achieve this by imposing movements up or down by a single grid point, imposing Step 2 above only with a certain probability chosen to deliver the correct average shock size.

A.4 Computing Historical Decompositions

In a classic linear setting, performing historical decompositions such as the one used in Section 6 is typically a trivial matter of inverting a data path using simple linear algebra. However, our

nonlinear model with heterogeneity and a discretized productivity process poses some additional computational challenges. Given a path for investment growth to match ΔI_t, $t = 1, ..., T^{decomp}$, we proceed as follows.

First, we pick an initial period drawn from a representative location in the unconditional simulation of the model with macro TFP equal to the steady-state level and the associated simulated cross-sectional distribution of firm-level states μ_0 drawn from the simulation of the model. Call this period $t = 0$, and note that at the end of period 0 a cross-sectional distribution μ_1 is pre-determined. Then for each period $t = 1, ..., T^{decomp}$, do the following:

1. Guess a value for macro TFP A_t, and find the bracketing interval $[A_{i-1}, A_i]$ together with linear interpolation weights $\omega(A_t, i) = \frac{A_t - A_{i-1}}{A_i - A_{i-1}}$ for the guessed value of productivity.

2. Compute the implied investment policies of all firms in the cross-sectional distribution μ_t given a macro TFP level equal to A_i, together with the implied macro investment level $I(A_i)$. Repeat the process for macro TFP equal to A_{i-1} to obtain $I(A_{i-1})$.

3. Assume that firms play a "mixed strategy" over the two macro TFP grid points, in which case the resulting macro investment level is $(1 - \omega(A_t, i))I(A_{i-1}) + \omega(A_t, i)I(A_i)$.

4. If the implied macro investment levels does not yield the desired investment growth value ΔI_t to within some tolerance, then update your guess for macro TFP A_t and return to Step 1. Otherwise proceed.

5. Given a productivity guess which delivers exactly the correct interpolated value of macro productivity in period t, compute the beginning-of-period distribution μ_{t+1} of firm-level states by pushing forward a fraction $\omega(A_t, i)$ of the distribution μ_t using firm policies associated with A_i and a fraction $1 - \omega(A_t, i)$ of the distribution μ_t using firm policies associated with A_{i-1}.

At the end of this process, you have determined a smooth value of productivity A_t which gives you an implied macro investment value exactly consistent with the target value in period t, and you have updated the cross-sectional distribution in an internally consistent fashion given the smooth value of productivity between grid points. Repeating this process for each period $t = 1, ..., T^{decomp}$ yields a productivity path A_t, as well as a set of cross-sectional distribution μ_t, which exactly match the target data path for investment. All other macro aggregates of interest can then be computed from the distributional and macro TFP path.

A.5 General Equilibrium Solution Algorithm

We follow an outer loop/inner loop approach to solving the model with endogenous wages.

1. Guess a mapping from a truncated history of aggregate states $(A_{t-K}, A_{t-K+1}, ..., A_{t-1}, A_t) \rightarrow W_t$.

2. Solve the model conditional upon this tractable truncated history, where $A_t, A_{t-1}, ..., A_{t-K}$ enter the firm's state vector and hence Bellman equations determining investment, default, and debt issuance policies.

3. Simulate the model for a large number of periods $t = 1, ..., T$, clearing markets with W_t in each period t by solving the nonlinear equation

$$\left(\frac{W_t}{\omega}\right)^{\lambda} = \int n(s, k, b|W_t) d\mu_t(s, k, b)$$

   for each period t in the simulation.

4. Based on the simulated wage data, update your wage prediction mapping from Step 1. If the mapping has converged, exit. If not update the mapping and return to Step 1.

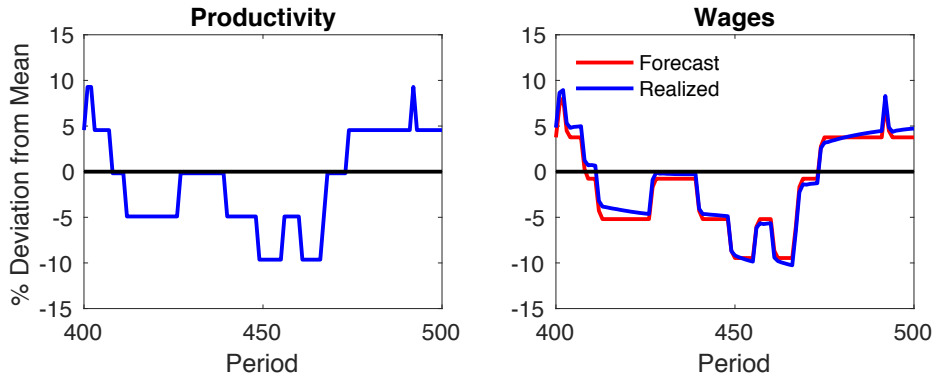


Figure 1: GE Predicted Wages

Notes: The left panel of the figure plots exogenous macro TFP realizations for a representative portion of the unconditional simulation of the baseline DE model. The right panel plots the corresponding wage realizations (blue) and forecast wages (red) according to the firm's stored vector of macro TFP histories.

This procedure works well in practice, delivering procyclical equilibrium wages close to their predicted values for our baseline choice of 1 lags of macro TFP. Figure 1 displays a representative portion of the unconditional simulation of the model, with macro TFP, equilibrium wages, and predicted wages.

A few practical comments are in order. First, given the discretized aggregate macro TFP state space, we store the wage mapping nonparametrically as a matrix of mean wages conditional upon each combination of truncated macro TFP histories. After simulation, the wage mapping update step simply involves repeated calculations of mean wages within the appropriate subsamples of

the simulated data. Second, because the aggregate state is replaced with macro TFP shock histories rather than with an augmented endogenous macro moment, there is no need to create an approximate anticipated default rule used to price debt. Lenders simply price debt according to the usual no-arbitrage conditions, conditional upon firm default policies which now have as explicit inputs the macro TFP shock histories. Third, because no endogenous moments are forecasted in our solution method, there is no Den Haan (2010)-style distinction between static and dynamic forecasts of the wage. In other words, there is no room for forecast errors about endogenous macro moments to accumulate over time, since only exogenous shock histories are used for forecasts. With this in mind, Figure 2 plots the estimated R^2 of regressions of the log wage on fully populated sets of dummies for macro TFP histories of up to a given lag length. Once a single lag is taken into account, incorporating information from yesterdays' TFP level about the current distribution of capital and hence labor demand in the cross section, the R^2 measures stabilize. Our baseline case with a single lag in the wage prediction rule is therefore a parsimonious but apparently accurate choice.

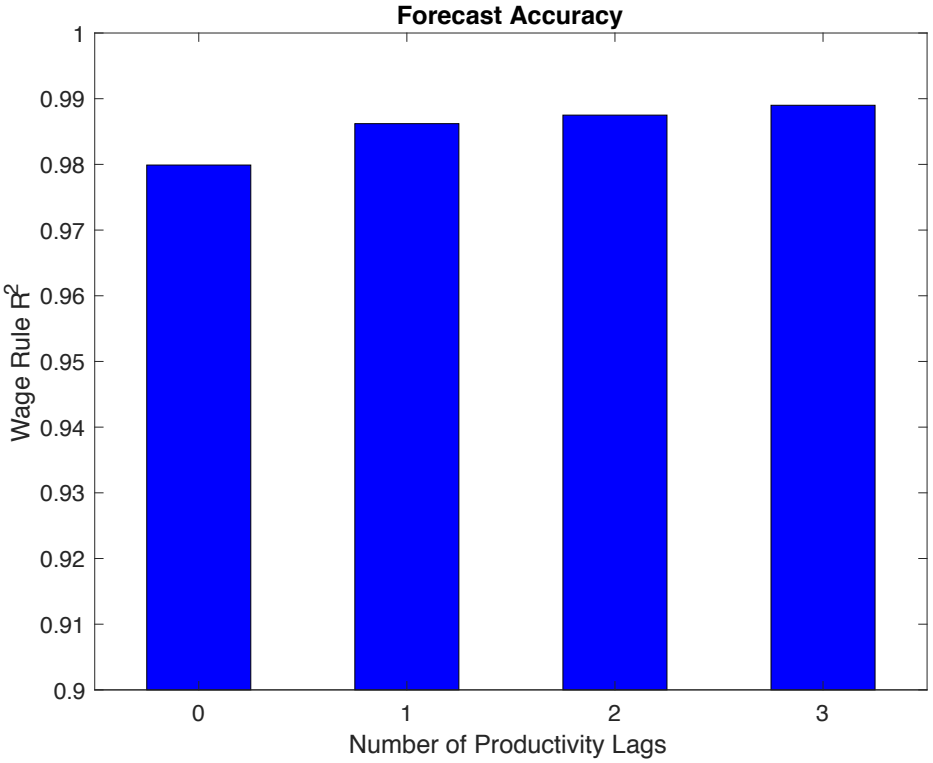


Figure 2: Wage Predictions and TFP Lags

Notes: The figure plots the R^2 of regressions of log wages on histories of macro TFP of increasing length.

B Data

B.1 Microdata on Firm Beliefs from Compustat and IBES-Guidance

We use a combination of the Compustat Fundamentals Annual and IBES manager guidance databases. The combined sample for the Compustat-IBES data in the calibration of the model spans 2007-2015 for 4457 firms. Descriptive statistics for each of the relevant variables used in moment construction, as well as firm revenues and capital book values, are reported in Table 2. Robustness checks for the forecast error reversion regression in Section 1 in the main text follow below.

Table 2: Sample Descriptive Statistics

Quantity	Mean	Standard Deviation
Sales	6200.09	23057.7
Capital	3997.236	10417.77
$\frac{\pi}{k}$	0.1617237	0.3406408
$\frac{i}{k}$	0.1859243	0.1124182
$\frac{b'}{k}$	1.394405	0.5345713
$\frac{fe'}{k}$	-0.0248892	0.2347275

Notes: The table reports descriptive statistics for the sample of 4457 firms from 2007-2015 in the combined Compustat-IBES database used to compute target moments. The first two rows represent revenues and the book value of the capital stock, in $ millions. The remaining rows reflect the ratio of realized earnings to the book value of the capital stock $\frac{\pi}{k}$, the capital expenditures investment rate $\frac{i}{k}$, the ratio of end of period total liabilities to the capital stock $\frac{b'}{k}$, and the next-period forecast error $\frac{fe'}{k}$, defined as realized future profits minus manager guidance scaled by firm capital. The sample was winsorized before computing the descriptive statistics above as well as computing the target moments.

The variable definitions are given as follows, with both empirical and model information attached:

- **Earnings** π are equal to GAAP net income, Compustat `ib`. The model equivalent is $\pi = (1-\tau)(y - Wn - AC(i,k)) + \tau(Rb + \delta k) - \delta k$.
- **Capital** k is equal to the book value of plants, property, and equipment, Compustat `ppent`. The model equivalent is the state variable k.
- **Investment** i is equal to the total value of capital expenditures, Compustat `capxv`. The model equivalent is the policy variable $i = k' - (1-\delta)k$.
- **Debt** b is equal to the total value of liabilities, Compustat `lt`. The model equivalent is the state variable b.

- **Forecast error** fe is equal to the realized value of earnings π minus the forecast level of earnings π^f made from the previous fiscal year, where realized earnings are Compustat `ib` and forecast earnings are equal to manager guidance from the IBES database. The model equivalent is the earnings value π above, minus the forecast level implied by firm-level diagnostic expectations, the definition of π, and firm policies predetermined in the previous period.

We also use the merged Compustat-IBES guidance sample to run various robustness checks to the firm forecast error predictability regressions reported in the main text. Table 3 shows similar forecast error predictability maintains after the Great Recession, and Table 4 shows that forecast error predictability is robust in a sample of firms present for five or more years in the data.

Table 3: Predictable Forecast Errors Post-Great Recession

	(1)	(2)	(3)	(4)
	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$
Profits$_t$	-0.133***			
	(0.030)			
Investment$_t$		-0.398***		
		(0.107)		
Debt Issuance$_t$			-0.466***	
			(0.147)	
Forecast$_t$				-0.320***
				(0.052)
Firm Effect	X	X	X	X
Time Effects	X	X	X	X
Years	2010-17	2010-17	2010-17	2010-17
Firm-Years	3880	3880	3880	3880

Notes: The table reports estimates of specifications on the merged Compustat - IBES Guidance sample at the firm-fiscal year level, restricting to the post-Great Recession period. Profits are earnings, investment is tangible capital expenditures, debt issuance is end-of-period debt, and forecast errors are actual earnings minus manager forecast earnings at a 1-period horizon. All series are relative to firm tangible capital stocks. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, and ***=1% level. The standard deviation of future forecast errors is 0.805, the standard deviation of profits is 1.037, the standard deviation of investment is 0.221, the standard deviation of debt issuance is 0.259, and the standard deviation of forecasts is 0.921. For all series, 0.01=1%.

Table 4: Predictable Forecast Errors: Firms Present for 5 or More Years

	(1)	(2)	(3)	(4)
	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$	Fcst. Error$_{t+1}$
Profits$_t$	-0.036*			
	(0.021)			
Investment$_t$		-0.451***		
		(0.066)		
Debt Issuance$_t$			-0.189**	
			(0.085)	
Forecast$_t$				-0.243***
				(0.030)
Firm Effect	X	X	X	X
Time Effects	X	X	X	X
Years	2010-17	2010-17	2010-17	2010-17
Firm-Years	9566	9566	9566	9566

Notes: The table reports estimates of specifications on the merged Compustat - IBES Guidance sample at the firm-fiscal year level, restricting to firms in the data for 5 or more years. Profits are earnings, investment is tangible capital expenditures, debt issuance is end-of-period debt, and forecast errors are actual earnings minus manager forecast earnings at a 1-period horizon. All series are relative to firm tangible capital stocks. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, and ***=1% level. The standard deviation of future forecast errors is 0.806, the standard deviation of profits is 1.051, the standard deviation of investment is 0.236, the standard deviation of debt issuance is 0.280, and the standard deviation of forecasts is 0.932. For all series, 0.01=1%.

B.2 Microdata on Bonds from FISD-TRACE

We use the WRDS US Corporate Bond Return database, which merges the Mergent FISD and FINRA TRACE datasets with issuance and secondary market information on corporate bond issues, respectively. We consider only unsecured, unconvertible debentures and convert secondary market yields to spreads based on comparable Treasury rates, with a resulting dataset of around 80,000 issues from mid-2002 to late 2019. We link the bond return database to Compustat firm financials through the WRDS CRSP link, and we aggregate from the issue to firm level by computing average yields and bond returns for a firm in Q4 of a given year. The resulting dataset spans around 1,500 large US public firms. Linking this panel to the IBES-manager guidance data yields the sample used in Table 2 in the main text. Table 5 replicates Table 2 conditioning only on investment grade bonds.

Table 5: Linking Forecast Errors and Credit Spreads: Investment Grade Only

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	Return_t	Return_t	Return_t	$\Delta\ \text{Spread}_t$	$\Delta\ \text{Spread}_t$	$\Delta\ \text{Spread}_t$
Estimation	OLS	IV	Red. Form	OLS	IV	Red. Form
Fcst. Error_t	0.001	0.083**		-0.001**	-0.085**	
	(0.001)	(0.042)		(0.001)	(0.043)	
Spread_{t-2}			0.185***			-0.190***
			(0.037)			(0.032)
First Stage		Fcst. Error_t			Fcst. Error_t	
Spread_{t-2}		2.235**			2.235**	
		(1.031)			(1.031)	
Years	2004-18	2004-18	2004-18	2004-18	2004-18	2004-18
Firm-Years	1691	1691	1691	1691	1691	1691
Firm Effects	X	X	X	X	X	X
Time Effects	X	X	X	X	X	X
First Stage F		4.7			4.7	

Notes: The table reports estimates of specifications on the merged Compustat - FISD/TRACE sample at the firm-fiscal year level. The top panel plots OLS, second-stage IV, and IV reduced-form estimates. The bottom panel, where relevant, reports first-stage estimates. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, and ***=1% level. The standard deviation of the bond return is 0.009, the standard deviation of spread growth is 0.015, the standard deviation of the forecast error is 0.373, and the standard deviation of the twice lagged spread level is 0.013. All series are in fractional units, so 0.01=1%.

B.3 Macro Data

At the macro level, we use a combination of information from the NIPA accounts, the Flow of Funds, and Moody's. The following variables are relevant, all at annual frequency or converted to annual frequency by averaging.

- **Output** Y is real GDP from the NIPA accounts in the data. In the model this is the total value of the firm-level outcome $y = Azk^{\alpha}n^{\nu}$ aggregated from the cross-sectional distribution.
- **Investment** I is real nonresidential private investment from the NIPA accounts in the data. In the model this is a choice variable for each firm, aggregated from the cross-sectional distribution.
- **Capital** K is real total private fixed assets from the NIPA accounts in the data. In the model this is the aggregated value of the state variable k from the cross-sectional distribution.
- **Spreads** are the Moody's BAA spread relative to 10-year Treasury bonds, at an annualized rate, in the data. In the model, the BAA spread is defined as the average of spreads in the

40-60th %-ile of spreads in the cross-sectional distribution of firms with positive spreads, loosely corresponding to Moody's definitions of the BAA spread.

- **Debt** B is total nonfinancial corporate debt from the Flow of Funds in the data. In the model this is the aggregated value of the state variable b from the cross-sectional distribution.
- **Macro Leverage** $\frac{B}{K}$ is the ratio of total debt to total capital, with each series defined as above.
- **Macro Investment Rate** $\frac{I}{K}$ is the ratio of total investment to total capital, with each series defined as above.

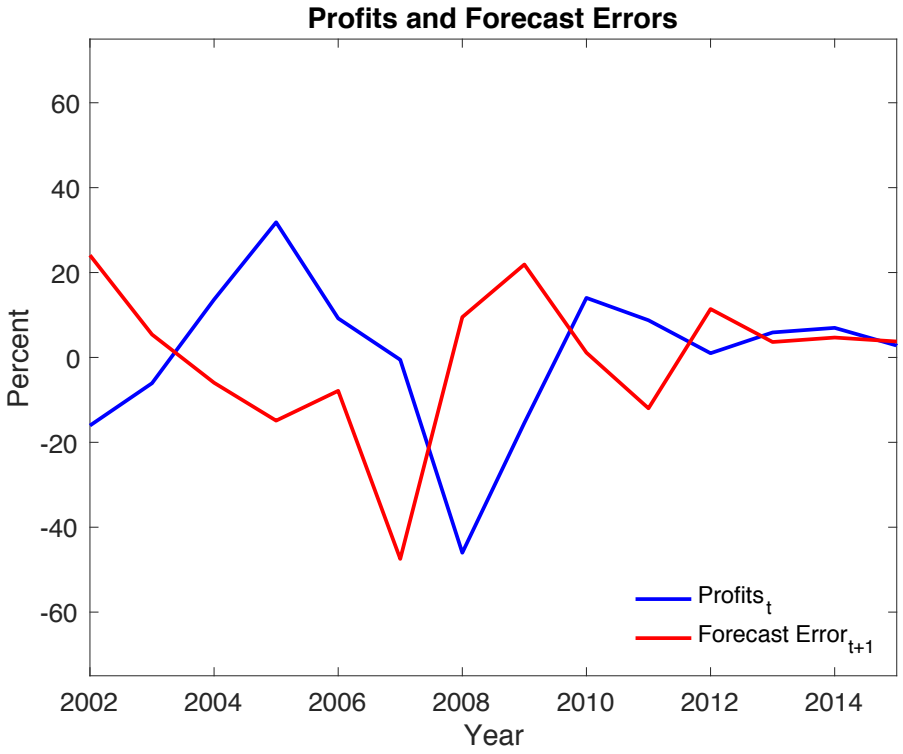
We use the macro series in multiple places. In the motivating Figure 1, we plot the empirical value of the Moody's BAA spread at quarterly frequency in recent years. In the empirical business cycle moments in Table 9, we reports moments from the growth rates (log differences of levels series and differences of percent series) of the indicated macro series over the common sample period 1986-2017. In the historical decomposition plots, we plot the dynamics of HP-filtered values of investment and debt, together with unadjusted spread growth or differences. Figure 3 plots the mean profit rates in year t against the mean profit forecast errors in $t+1$ from the Compustat-IBES Guidance dataset. Table 6 reports covariances of the mean profit rate in t with other macro aggregates computes from the Compustate-IBES Guidance dataset and relating to firm investment, debt issuance, and future forecast errors.

Table 6: Predictable Reversals in Macro Forecast Errors

	Profits_t	Investment_t	Debt Issuance_t	Fcst. Error_{t+1}
Profits_t	1.00	0.32	0.63	-0.38
Investment_t	0.32	1.00	0.52	-0.67
Debt Issuance_t	0.63	0.52	1.00	-0.30
Fcst. Error_{t+1}	-0.38	-0.67	-0.30	1.00
Std. Deviation	0.178	0.046	0.067	0.163

Notes: The table reports correlations of mean profits, investment, debt issuance, and future forecast errors. In year t these series average across firms in the merged Compustat-IBES Guidance sample. All series are normalized by firm-level tangible capital stocks. Forecast errors are realized minus actual from manager guidance of earnings. Sample 2002-16.

Figure 3: Predictable Reversals in Macro Forecast Errors



Notes: The figure plots the yearly mean profits across firms in the merged IBES Guidance - Compustat database in blue. The red line in a given year plots the mean forecast error, actual minus forecast, for the next year, averaging across firms. Both measures are expressed as percentages of firm-level tangible capital stocks and demeaned. A time series regression of the future forecast error on profits today yields an estimated slope coefficient of -0.42 (p-value $= 0.002$) with 5 lags for Newey-West standard errors over the 2002-2016 period.