The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. , NO. 0 • XXXX 2017

Diagnostic Expectations and Credit Cycles

PEDRO BORDALO, NICOLA GENNAIOLI, and ANDREI SHLEIFER*

ABSTRACT

We present a model of credit cycles arising from diagnostic expectations—a belief formation mechanism based on Kahneman and Tversky's representativeness heuristic. Diagnostic expectations overweight future outcomes that become more likely in light of incoming data. The expectations formation rule is forward looking and depends on the underlying stochastic process, and thus is immune to the Lucas critique. Diagnostic expectations reconcile extrapolation and neglect of risk in a unified framework. In our model, credit spreads are excessively volatile, overreact to news, and are subject to predictable reversals. These dynamics can account for several features of credit cycles and macroeconomic volatility.

THE FINANCIAL CRISIS OF 2008 TO 2009 REVIVED INTEREST among economists and policy makers in the relationship between credit expansion and subsequent financial and economic busts. According to an old argument (e.g., Minsky (1977)), investor optimism brings about the expansion of credit and investment, and leads to a crisis when such optimism abates. Stein (2014) echoes this view by arguing that policy makers should be mindful of credit market frothiness and consider countering it with policy. In this paper, we develop a behavioral model of credit cycles with microfounded expectations that is consistent both with the Minsky narrative and with a great deal of evidence.

Recent empirical research has documented a number of credit cycle facts. Using a sample of 14 developed countries between 1870 and 2008, Schularick and Taylor (2012) demonstrate that rapid credit expansions forecast declines in real activity. Jorda, Schularick, and Taylor (2013) further find that more credit-intensive expansions are followed by deeper recessions. Mian, Sufi, and Verner (2017) show that growth in household debt predicts economic slowdowns, Baron and Xiong (2017) show for a sample of 20 developed countries that bank credit expansion predicts increased crash risk in both bank stocks and equity markets

*Pedro Bordalo is at Saïd Business School, University of Oxford. Nicola Gennaioli is at Universita Bocconi and IGIER. Andrei Shleifer is at Harvard University. Gennaioli thanks the European Research Council and Shleifer thanks the Pershing Square Venture Fund for Research on the Foundations of Human Behavior for financial support of this research. The authors are also grateful to Nicholas Barberis, Bruce Carlin, Lars Hansen, Sam Hanson, Arvind Krishnamurthy, Gordon Liao, Yueran Ma, Matteo Maggiori, Sendhil Mullainathan, Andreas Schaab, Josh Schwartzstein, Jesse Shapiro, Alp Simsek, Jeremy Stein, Amir Sufi, David Thesmar, Chari Varadarajan, Wei Xiong, Luigi Zingales, and two anonymous referees for helpful comments. The authors have no conflicts of interest to declare.

DOI: 10.1111/jofi.12586

more broadly, and Fahlenbrach, Prilmeier, and Stulz (2016) find, in a cross-section of U.S. banks, that rapid loan growth predicts poor loan performance and low bank returns in the future.

Similar findings emerge from an examination of credit conditions. Greenwood and Hanson (2013) show that the credit quality of corporate debt issuers deteriorates during credit booms, and that a high share of risky loans forecasts low, and even negative, corporate bond returns. Gilchrist and Zakrajšek (2012) and Krishnamurthy and Muir (2015) show that credit tightening correctly anticipates recessions. Lopez-Salido, Stein, and Zakrajsek (2017) find that low credit spreads predict both a rise in credit spreads and low economic growth afterwards. They stress predictable mean-reversion in credit market conditions.[1] In Section I, we offer preliminary evidence that survey forecasts of credit spreads are excessively optimistic when these spreads are low, and that both errors and revisions in forecasts are predictable. Overall, the existing evidence is hard to square with rational expectations, indicating a need for a behavioral approach to modeling credit cycles.

We propose a behavioral model that both accounts for this evidence and describes in a dynamic setup how credit markets overheat. We begin with a psychologically founded model of beliefs and their evolution in light of new data.[2] This model was developed to account for judgment biases that are well documented in the lab, such as the conjunction and disjunction fallacies and base rate neglect, and is therefore portable in the sense of Rabin (2013). The model relies on Gennaioli and Shleifer's (2010) formalization of Kahneman and Tversky's (1972) representativeness heuristic. According to Kahneman and Tversky, a certain attribute is judged to be excessively common in a population when that attribute is diagnostic for the population, meaning that it occurs more frequently in the given population than in a relevant reference population (Tversky and Kahneman (1983)). For example, after seeing a patient test positive on a medical test, doctors tend to overestimate the likelihood that the patient has the disease because sick people are more frequent in the population of positive tests relative to the population of negative tests, even when they are few in absolute terms (Casscells, Schoenberger, and Graboys (1978)).

This idea can be naturally applied to modeling expectations in a macroeconomic context. Similar to the medical test example, agents overweight those future states whose likelihood increases the most in light of current news *relative to what they know already*. Thus, just as doctors overestimate the probability of sickness after a positive test result, agents overestimate the probability of a

[1] See also Bernanke (1990), Friedman and Kuttner (1992), and Stock and Watson (2003), among others.

[2] Many models of beliefs in finance are motivated by psychological evidence, but often use specifications specialized to financial markets (e.g., Muth (1961), Barberis, Shleifer, and Vishny (1998), Rabin and Vayanos (2010), Fuster, Laibson, and Mendel (2010), Hirshleifer, Li, and Yu (2015), Greenwood and Hanson (2015), Barberis et al. (2015)). Fuster, Laibson, and Mendel (2010) review lab and field evidence on deviations from rational expectations.

good future state when the current news is good. Because agents overweight diagnostic information, we refer to such beliefs as diagnostic expectations.

The above approach has significant implications. A path of improving news leads an agent to focus on good future outcomes and neglect the bad ones, causing excessive optimism, while a path of deteriorating news leads the agent to focus on bad future outcomes and neglect good ones, causing excessive pessimism. There is a kernel of truth in assessments: revisions respond to news, but excessively. When change slows down there is a reversal, so that crises can occur even in good times without deteriorating fundamentals. Expectations are excessively volatile.

Our model unifies the phenomena of extrapolation (Cagan (1956), Cutler, Poterba, and Summers (1990), DeLong et al. (1990), Barberis and Shleifer (2003), Greenwood and Shleifer (2014), Barberis et al. (2015, 2016), Gennaioli, Ma, and Shleifer (2016)) and the neglect of risk (Gennaioli, Shleifer, and Vishny (2012), Coval, Pan, and Stafford (2014), Arnold, Schuette, and Wagner (2015)). Critically, under diagnostic expectations, departures from the rational benchmark are driven by updates in the probability of future events, and thus depend on the true distribution of states of nature. Unlike in mechanical extrapolation, distortions arise only when news is informative about future events. The model converges to rational expectations when the process is not persistent. On the other hand, distortions such as neglect of tail risk depend on the true volatility of the process. In particular, neglect of tail risk is stronger when volatility is low. In this case, news is actually more informative about changes in the probability of future states.

We construct a neoclassical macroeconomic model in which the only nonstandard feature is expectations. To begin, we do not include financial or any other frictions. The model accounts for many empirical findings, some of which also obtain under rational expectations, but some of which do not. In our model:

Prediction 1: In response to good news about the economy, credit spreads decline, credit expands, the share of high-risk debt rises, and investment and output grow.

Prediction 2: Following the period of narrow credit spreads, these spreads rise on average, credit and the share of high-risk debt decline, and investment and output decline as well. Larger spikes in spreads predict lower GDP growth.

Prediction 3: Credit spreads are too volatile relative to fundamentals and their changes are predictable in a way that parallels the cycles described in Predictions 1 and 2.

Prediction 4: Investors make predictable forecast errors and forecast revisions. Bond returns are also predictable in a way that parallels Predictions 1 and 2.

Prediction 1 can obtain under rational expectations, and the same is true about Prediction 2 provided fundamentals are mean-reverting. Predictions 3 and 4, in contrast, critically depend on our model of diagnostic expectations. After presenting the basic model, we briefly explore the interaction between diagnostic expectations and leverage by proposing an extension with a preferred habitat for safe debt, a version of financial frictions. In this setting, excess volatility in expectations causes strong market reactions by inducing a reclassification of debt.

Our paper is related to three strands of research. First, the prevailing approach to relating financial markets and the real economy is through the use of financial frictions, which focus on the transmission of an adverse shock through a leveraged economy (Bernanke and Gertler (1989), Kiyotaki and Moore (1997)).[3] The adverse shock in such models is either a drop in fundamentals or a "financial shock" consisting of a tightening of collateral constraints or an increase in required returns. However, these models do not usually explain the sources of such shocks. Moreover, because they assume rational expectations, these models do not explain predictable negative or low abnormal returns on debt in overheated markets nor do they explain systematic errors in expectations. Our model accounts for both market crises and abnormal returns.

Second, our paper is related to recent work on limited attention (e.g., Sims (2003), Gabaix (2014)). In general, these models predict sluggish expectations and underreaction to information, consistent with empirical evidence for inflation (Coibion and Gorodnichenko (2012, 2015)). Also related is research on momentum and slow reaction to information in financial markets (Jegadeesh and Titman (1993), Hong and Stein (1999), Bouchaud et al. (2016)). Our model most naturally delivers overreaction to information, although we discuss briefly how the two approaches can be unified.

Third, our paper extends research on behavioral credit cycles initiated by Minsky (1977). Gennaioli, Shleifer, and Vishny (2012) present a model of neglected risk. Gennaioli, Shleifer, and Vishny (2015) sketch a model of credit cycles exhibiting both under- and overreaction based on the Bordalo et al. (2016) model of stereotypes. Jin (2015) models extrapolation in credit markets. Greenwood, Hanson, and Jin (2017) present a model of extrapolation of default rates that also delivers many credit cycle facts.

In Section I, we document the predictability of both forecast errors and forecast revisions. Section II introduces diagnostic expectations, describes how they evolve, and relates our formulation to extrapolation and neglect of risk. Section III presents our model of credit cycles. Section IV develops the predictions of the model for the behavior of credit spreads, expectations about credit spreads, and the link between spreads and economic activity. Section V incorporates safe debt. Section VI concludes. An Internet Appendix contains all proofs and discusses some alternative specifications of the diagnostic expectations model.[4]

[3] Some papers add Keynesian elements to financial frictions, such as the zero lower bound on interest rates or aggregate demand effects (e.g., Eggertsson and Krugman (2012), Rognlie, Shleifer, and Simsek (2017)).

[4] The Internet Appendix is available in the online version of this article on *The Journal of Finance's* website.

I. Evidence on Expectations and Credit Spreads

We start by offering some motivating evidence on analysts' expectations regarding the Baa bond – Treasury credit spread, a commonly used indicator of credit market conditions (Greenwood and Hanson (2013)). We do not have enough data to perform a stringent test of rational expectations, but we can illustrate how an analysis of credit cycles can be deepened with expectations data, and establish some facts that a model of expectations formation should account for.

We use data from Blue Chip Financial Forecasts, a monthly survey of around 40 panelists from major financial institutions.[5] They provide forecasts of various interest rates for the current quarter t and for quarters $t + 1$ through $t + 4$. We use the consensus forecast, which is the mean across analysts. We take data from the March, June, September, and December publications. Forecasts of the Baa spread are obtained as the difference between the forecasts of the Baa corporate bond yield and the 10-year Treasury yield. Our data cover the period 1999Q1 to 2014Q4.

A. *Predictability in Forecast Errors*

Under the assumption of rational expectations (and knowledge of the data generating process), analysts' forecast errors should be unpredictable from past data. Figure 1 plots, over time, the current spread (averaged over quarters $t - 4$ to $t - 1$) against forecast errors of the credit spread over the next year (averaged over quarters $t + 1$ to $t + 4$). The data suggest predictability: when the current spread is low, the expected future spread is too low, and when the current spread is high, the expected future spread is too high. The 1999 to 2000 and 2005 to 2008 periods are associated with low spreads and excessive optimism, while the early 2000s and the recent crisis are associated with high spreads and excessive pessimism.

Table I reports results for an econometric test of predictability. Column (1) presents results for the regression of the actual credit spread (averaged over the next four quarters) on the current spread (averaged over the past four quarters), column (2) for the corresponding analyst forecast regressed on the current spread, and column (3) for the forecast error (actual minus forecast) regressed on the current spread. Table I confirms the message of Figure 1. In column (3), the higher the current spread, the higher the forecast relative to the realization. This may occur because analysts perceive excessive persistence in current conditions: in column (1), the estimated persistence of the credit spread over the next four quarters is about 0.4, while in column (2), the forecasted credit spread implies a coefficient of about 0.6.

[5] One may worry that Blue Chip professional forecasts are distorted for signaling or entertainment reasons since participants are not anonymous. However, these forecasts tend to be very similar to the anonymous forecasts collected by the Philadelphia Fed Survey of Private Forecasters. Moreover, unlike in the case of stock analysts, there is no unconditional bias in the Blue Chip forecasts we study here.

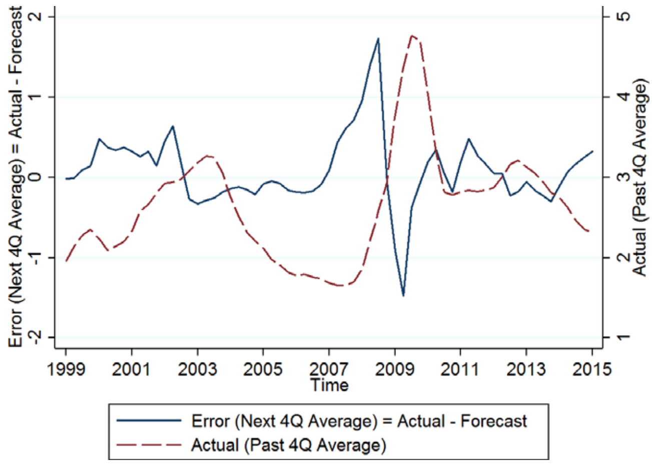

Figure 1. Predictable errors in forecasts of credit spreads. Quarterly time series plot: in each quarter t, the solid line shows errors (actual minus forecast) associated with the contemporaneous forecasts of the credit spread, averaged over quarters $t+1$ to $t+4$ (left scale), and the dashed line shows the credit spread averaged over quarters $t-4$ to $t-1$, where $t-1$ is the latest quarterly credit spread prior to the forecast (right scale). Credit forecasts are the consensus forecasts computed from Blue Chip Financial Forecasts surveys. (Color figure can be viewed at wileyonlinelibrary.com)

Table I

Actual, Forecast, and Error of Future Credit Spreads

Quarterly time series regressions: the dependent variable is the actual credit spread averaged over quarters $t+1$ to $t+4$ in column (1), quarter t forecasts of credit spreads averaged over quarters $t+1$ to $t+4$ in column (2), and the forecast error (actual−forecast) of credit spreads in column (3); the independent variable is the actual credit spread averaged over quarters $t-4$ to $t-1$, where $t-1$ is the latest quarterly credit spread prior to the forecast. Standard errors in parentheses are Newey-West, with the automatic bandwidth selection of Newey and West (1994).

	Actual (1)	Forecast (2)	Error (3)
Current spread	0.39	0.65	−0.26
	(1.67)	(4.62)	(−2.20)
Constant	1.63	0.86	0.77
	(2.56)	(2.25)	(2.40)
Observations	64	64	64
R^2	0.16	0.47	0.16

B. Tests of Expectations' Revisions

We next examine forecast revisions, which should also be unpredictable under rational expectations. Figure 2 plots the current spread against the future forecast revision, defined as the difference between the forecast for the spread in quarter $t + 4$ made in quarter $t + 3$ and the current (quarter t) forecast for the same spread. The evidence, confirmed in the econometric test of Table II,

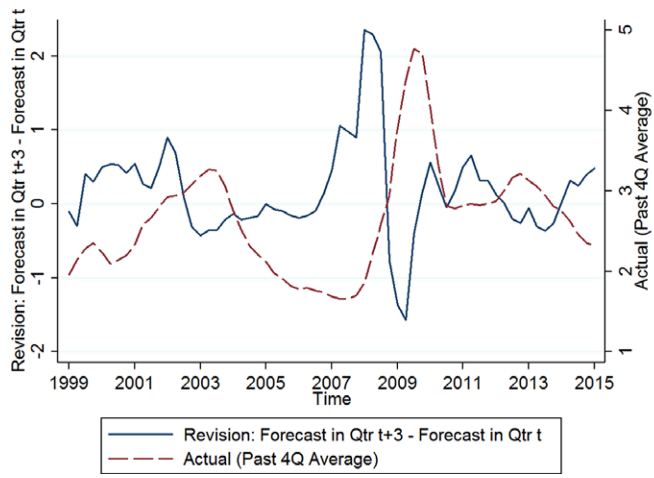

Figure 2. Predictable revisions in forecasts of credit spreads. Quarterly time series plot: in each quarter t, the solid line shows forecast revisions (quarter $t+3$ forecast of credit spread in quarter $t+4$ minus quarter t forecast of credit spread in quarter $t+4$, left scale), and the dashed line shows the credit spread averaged over quarters $t-4$ to $t-1$, where $t-1$ is the latest quarterly credit spread prior to the forecast (right scale). Credit forecasts are the consensus forecast computed from Blue Chip Financial Forecasts surveys. (Color figure can be viewed at wileyonlinelibrary.com)

Table II

Forecast Revisions of Credit Spreads

Quarterly time series regression: the dependent variable is the forecast revision (quarter $t+3$ forecast of credit spread in quarter $t+4$ minus quarter t forecast of credit spread in quarter $t+4$); the independent variable is actual credit spread averaged over quarters $t-4$ to $t-1$, where $t-1$ is the latest quarterly credit spread prior to the forecast. Standard errors in parentheses are Newey-West, with the automatic bandwidth selection of Newey and West (1994).

	Revision (1)
Current spread	−0.36
	(−2.13)
Constant	1.13
	(2.44)
Observations	64
R^2	0.15

again suggests predictability: when the current spread is low, forecasts are revised upwards, whereas when the current spread is high, forecasts are revised downwards.

This evidence is difficult to reconcile with rational expectations, but suggests that analysts' forecasts follow a boom-bust pattern. During booming bond markets (low spreads), expectations are too optimistic and systematically revert in the future, planting the seeds for a cooling of bond markets. The extrapolative dynamics of forecasts that we document here are in line with other studies, such as evidence of systematic reversal in bond spreads (Greenwood and Hanson (2013)) and the extrapolative nature of a CFO's expectations about their company's earnings growth (Gennaioli, Ma, and Shleifer (2016)).

II. Diagnostic Expectations

A. Formal Model of Representativeness

We build our model of expectations from first principles, starting with Kahneman and Tversky's representativeness heuristic, which they define as follows: "an attribute is representative of a class if it is very diagnostic; that is, the relative frequency of this attribute is much higher in that class than in the relevant reference class" (Tversky and Kahneman (1983, p. 296)). They argue that individuals often assess likelihood by representativeness, thus estimating types or attributes as being likely when they are instead representative, and present a great deal of experimental evidence to support this claim. Gennaioli and Shleifer (2010) construct a model in which judgment biases arise because decision makers overweight events that are representative precisely in the sense of Kahneman and Tversky's definition. To motivate diagnostic expectations, we briefly describe this model and a related application to stereotype formation by Bordalo et al. (2016).

A decision maker judges the distribution of trait T in group G. The true distribution of the trait is $h(T = t|G)$. Gennaioli and Shleifer (2010) define the representativeness of trait $T = t$ for group G as

$$\frac{h(T = t|G)}{h(T = t|-G)},$$

where $-G$ is a relevant comparison group. As in the quote above, a trait is more representative if it is *relatively* more frequent in G than in $-G$. Gennaioli and Shleifer (2010) assume that representative types are easier to recall. Due to limited working memory, an agent overweighs these types in his assessment. Note that in this model beliefs about group G are context dependent because they depend on features of the comparison group $-G$.

To illustrate, consider an individual assessing the distribution of hair color among the Irish. The trait T is hair color, and the conditioning group G is the Irish. The comparison group $-G$ is the world at large. The true relevant distributions are:[6]

	$T = red$	$T = blond/light\ brown$	$T = dark\ brown$
$G \equiv Irish$	10%	40%	50%
$-G \equiv World$	1%	14%	85%

[6] Shares of red hair are accurate; shares of light and dark brown are approximate. See http://www.eupedia.com/genetics/origins_of_red_hair.shtml and http://www.eupedia.com/europe/genetic_maps_of_europe.shtml.

The most representative hair color for the Irish is red because it is associated with the highest likelihood ratio among hair colors:

$$\frac{Pr\,(red\ hair\ |Irish)}{Pr\,(red\ hair\ |World)} = \frac{10\%}{1\%} = 10.$$

Our model thus predicts that assessments exaggerate the frequency of red-haired Irish. After hearing the news "Irish," the representative red-haired type comes to mind and its likelihood is inflated. The agent discounts the probability of blond and dark hair types because they are less available when thinking about the Irish.

In sum, representativeness causes an agent to inflate the likelihood of types whose *objective* probability rises the most in G relative to the reference *context* $-G$.[7] Gennaioli and Shleifer (2010) and Bordalo et al. (2016) show that this model unifies several widely documented errors in probabilistic judgment, such as base rate neglect (the medical test example in the introduction) and the conjunction fallacy, and sheds light on key features of social stereotypes and context-dependent beliefs.

B. Diagnostic Expectations

The same logic can be applied to belief formation about aggregate economic conditions. Let time be discrete $t = 0, 1, \ldots$ The state of the economy at t is a random variable ω_t that follows the AR(1) process $\omega_t = b\omega_{t-1} + \epsilon_t$, with $\epsilon_t \sim N(0, \sigma^2)$, $b \in [0, 1]$. The model can be easily generalized to richer AR(N)-normal processes.

When forming a forecast the agent assesses the distribution of a certain future state, say $\hat{\omega}_{t+1}$, entailed by the realized current conditions $\omega_t = \hat{\omega}_t$. This is similar to the medical test example, where the doctor assesses the health of the patient conditional on a positive test outcome. Pursuing the analogy, the agent must predict the distribution of future prospects ω_{t+1} in a group $G \equiv \{\omega_t = \hat{\omega}_t\}$ that summarizes current conditions.

The rational agent solves this problem by using the true conditional distribution of ω_{t+1} given $\omega_t = \hat{\omega}_t$, denoted $h(\omega_{t+1}|\omega_t = \hat{\omega}_t)$. An agent whose judgments are shaped by representativeness has this true distribution in the back of his mind, but selectively retrieves and thus overweighs realizations of ω_{t+1} that are representative or diagnostic of $G \equiv \{\omega_t = \hat{\omega}_t\}$ relative to the background context $-G$. But what is $-G$ here?

In the Irish example, representativeness assesses $G = Irish$ against $-G =$ *Rest of the world*. In the medical test example, $G =$ *Positive test* is assessed against $-G =$ *Not taking the test*, so that context captures an absence of new

[7] This example also illustrates context dependence of beliefs. It is the scarcity of red-haired people in the "rest of the world" that renders red hair so distinctive for the Irish. The judgment bias would be less pronounced if the share of red-haired people in the rest of the world were to rise, or equivalently if the agent was primed to think about the Irish in the context of a more similar group (e.g., $-G = Scots$).

information. We adopt this dynamic perspective by taking context at time t as reflecting information held at $t-1$. Formally, context is the state prevailing if there is no news, which is $-G \equiv \{\omega_t = b\hat{\omega}_{t-1}\}$ under the assumed AR(1). A certain future state $\hat{\omega}_{t+1}$ is thus more representative at t if it is more likely to occur under the realized state $G \equiv \{\omega_t = \hat{\omega}_t\}$ than on the basis of past information $-G \equiv \{\omega_t = b\hat{\omega}_{t-1}\}$. Representativeness of $\hat{\omega}_{t+1}$ is then given by

$$\frac{h\left(\hat{\omega}_{t+1} \mid \omega_t = \hat{\omega}_t\right)}{h\left(\hat{\omega}_{t+1} \mid \omega_t = b\hat{\omega}_{t-1}\right)}. \tag{1}$$

The most representative state is the one exhibiting the largest increase in its likelihood based on recent news. The comparison group $-G$ could alternatively be slow moving, including more remote recollections. We discuss these cases in the Internet Appendix.[8]

The psychology of diagnostic expectations works as follows. Decision makers have in the back of their mind the conditional distribution $h(\hat{\omega}_{t+1}|\omega_t = \hat{\omega}_t)$. After seeing current news $\Omega_t = \omega_t$, those states whose probability increased the most come immediately to mind. Memory limits then imply that the agent oversamples these states from his memory database $h(\hat{\omega}_{t+1}|\omega_t = \hat{\omega}_t)$. Beliefs inflate the probability of more representative states and deflate the probability of less representative states.

We formalize overweighting of representative states "as if" the agent uses the distorted density

$$h_t^{\theta}\left(\hat{\omega}_{t+1}\right) = h\left(\hat{\omega}_{t+1} \mid \omega_t = \hat{\omega}_t\right) \cdot \left[\frac{h\left(\hat{\omega}_{t+1} \mid \omega_t = \hat{\omega}_t\right)}{h\left(\hat{\omega}_{t+1} \mid \omega_t = b\hat{\omega}_{t-1}\right)}\right]^{\theta} \frac{1}{Z},$$

where the normalizing constant Z ensures that $h_t^{\theta}(\hat{\omega}_{t+1})$ integrates to one, and $\theta \in [0, +\infty)$ measures the severity of judging by representativeness. When $\theta = 0$, the agent has no memory limits and appropriately uses all information, forming rational expectations. When $\theta > 0$, memory is limited. The distribution $h_t^{\theta}(\hat{\omega}_{t+1})$ then inflates the likelihood of representative states that come to mind quickly, while it deflates the likelihood of nonrepresentative states. Because they overweight the most diagnostic future outcomes, we call the expectations formed in light of $h_t^{\theta}(\hat{\omega}_{t+1})$ diagnostic.

In our analysis, we take θ to be fixed. In principle, however, recall can depend on the agent's deliberate effort, in which case θ may vary across situations. In Section IV, we calibrate θ given a model of fundamentals and the facts in Section I.

The distorted distribution $h_t^{\theta}(\hat{\omega}_{t+1})$ shows that news not only alters the objective likelihood of certain states but also changes the extent to which the agent focuses on them. An event that increases the likelihood of a future state $\hat{\omega}_{t+1}$

[8] The analysis is less tractable if context is defined to be the past state $-G \equiv \{\omega_{t-1} = \hat{\omega}_{t-1}\}$. In fact, the distributions $h(\hat{\omega}_{t+1}|\omega_t = \hat{\omega}_t)$ and $h(\hat{\omega}_{t+1}|\omega_{t-1} = \hat{\omega}_{t-1})$ have different variances, which distorts not only the mean but also the variance of the target distribution. The representation obtained in (2) below—and the tractability it entails—extends to other distributions, including lognormal and exponential.

also makes it more representative, so $h_t^\theta(\hat{\omega}_{t+1})$ overshoots. The reverse occurs when the likelihood of $\hat{\omega}_{t+1}$ decreases. If the likelihood ratio in (1) is monotone increasing, "rationally" good news causes overweighting of high future states and underweighting of low future states (the converse is true if news is bad). In this sense, good news causes neglect of downside risk. Diagnostic beliefs have a convenient representation.

PROPOSITION 1: *When the process for ω_t is AR(1) with normal $(0, \sigma^2)$ shocks, the diagnostic distribution $h_t^\theta(\hat{\omega}_{t+1})$ is also normal, with variance σ^2 and mean*

$$\mathbb{E}_t^\theta(\omega_{t+1}) = \mathbb{E}_t(\omega_{t+1}) + \theta\left[\mathbb{E}_t(\omega_{t+1}) - \mathbb{E}_{t-1}(\omega_{t+1})\right]. \tag{2}$$

See the Internet Appendix for all proofs. Proposition 1 is a representation result for diagnostic expectations in terms of the distributions held in the agent's memory. Diagnostic beliefs can be represented as a linear combination of the rational expectations of ω_{t+1} held at t and at $t-1$. If the agent had in mind a subjective distribution, based for instance on his personal knowledge, diagnostic expectations would entail a similar transformation of that distribution. The general point of equation (2) is the "kernel of truth" logic: diagnostic expectations overreact to the information received at t by the term $\theta[\mathbb{E}_t(\omega_{t+1}) - \mathbb{E}_{t-1}(\omega_{t+1})]$. Crucially, it is *not* the case that decision makers compute and combine rational expectations. Rather, oversampling representative future states of a specific random variable, as defined in (1), implies that the "news" term in the right-hand side of (2) excessively influences the agent's subjective beliefs. This is consistent with Kahneman's (2011, p. 324) view that "our mind has a useful capability to focus spontaneously on whatever is odd, different, or unusual." Figure 3 illustrates such neglect of risk. After good news, the diagnostic distribution of ω_{t+1} is a right shift of the objective distribution, which underestimates probabilities in the left tail (the shaded area). There are two main differences with the model of neglected risk in Gennaioli, Shleifer, and Vishny (2012), who assume that investors neglect unlikely events. First, under diagnostic expectations, neglect of tail events is selective: one tail is neglected, while the other tail is exaggerated. This is due to the emphasis on probability changes, not levels. After good news, investors neglect the left tail but overweight the right tail, even if the latter remains unlikely in absolute terms.

Second, neglect of risk depends on the volatility of economic conditions. Indeed, representativeness varies across future states of the world ω_{t+1} as follows:

$$\frac{\partial}{\partial\hat{\omega}_{t+1}} \ln\left[\frac{h(\hat{\omega}_{t+1}|\omega_t = \hat{\omega}_t)}{h(\hat{\omega}_{t+1}|\omega_t = b\hat{\omega}_{t-1})}\right] = \frac{\mathbb{E}_t(\omega_{t+1}) - \mathbb{E}_{t-1}(\omega_{t+1})}{\sigma^2}.$$

After good news, when $\mathbb{E}_t(\omega_{t+1}) > \mathbb{E}_{t-1}(\omega_{t+1})$, the above expression is positive and the least representative states are the lowest ones, particularly if volatility σ^2 is low. After bad news, bad states are the most representative ones, particularly when volatility σ^2 is low. As a result, a given diagnostic distortion of

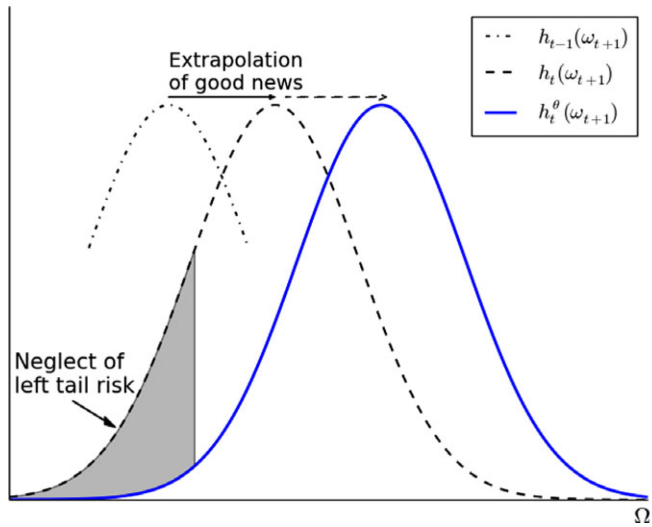

Figure 3. Neglect of risk and extrapolation. This plot shows the subjective distribution of fundamentals in period $t+1$, $h_t^\theta(\omega_{t+1})$, based on diagnostic expectations (solid line) after good news at t, the objective distribution of fundamentals in period $t+1$, $h_t(\omega_{t+1})$, based on rational expectations (dashed line), and the lagged distribution of fundamentals in period $t+1$, $h_{t-1}(\omega_{t+1})$, based on rational expectations (dash-dot line). (Color figure can be viewed at wileyonlinelibrary.com)

mean beliefs $\mathbb{E}_t^\theta(\omega_{t+1})$ causes larger distortions in the perception of tail risk under stable economic conditions.

This mechanism differentiates distortions of mean beliefs and distortions of tail events. In a stable environment, that is, when σ^2 is low, mild positive news causes a mild exaggeration of mean optimism $\mathbb{E}_t^\theta(\omega_{t+1})$ but a drastic neglect of downside risk. For example, suppose left-tail risk is defined as a realization $\omega_t < -2\sigma$. Then an increase in optimism due to diagnostic expectations, say from $\mathbb{E}_t(\omega_{t+1}) = 0$ to $\mathbb{E}_t^\theta(\omega_{t+1}) = \sigma$, entails a drop in perceived left-tail risk from 2.275% to 0.135%. The lower is volatility σ, the smaller is the improvement in the mean necessary to reduce the perception of tail risk. This aspect plays an important role for credit spreads, which reflect default risk and thus depend (even in the rational benchmark) on the volatility of economic conditions.[9]

One important consequence of selective neglect of tail events is that it connects neglect of tail risk with extrapolation by the same psychological mechanism. For the AR(1) process $\omega_t = b\omega_{t-1} + \epsilon_t$, with persistence parameter b, equation (2) becomes

$$\mathbb{E}_t^\theta(\omega_{t+1}) - \omega_t = \left[\mathbb{E}_t(\omega_{t+1}) - \omega_t\right] + b \cdot \theta \cdot \left[\omega_t - \mathbb{E}_{t-1}(\omega_t)\right],$$

[9] As we show in the Internet Appendix, Proposition 1 extends to GARCH processes in which not only the mean but also the variance varies over time. In this more general case, both the perceived mean and variance are distorted by news, and the distortion in mean beliefs $\mathbb{E}_t^\theta(\omega_{t+1})$ depends on the variance innovation.

that is, the current shock $\omega_t - \mathbb{E}_{t-1}(\omega_t)$ is extrapolated into the future, but only if the data are serially correlated, $b > 0$. The key difference with mechanical extrapolative formulas is that, under diagnostic expectations, the extent of extrapolation depends on the true persistence b of the process that the agent tries to forecast. This is a testable prediction that also makes our model immune to the Lucas critique.[10]

It is straightforward to extend diagnostic expectations to longer term forecasts.

COROLLARY 1: *When the process for ω_t is AR(1) with normal $(0, \sigma^2)$ shocks, the diagnostic expectation for ω_{t+T} is given by*

$$\mathbb{E}_t^\theta\left(\omega_{t+T}\right) = \mathbb{E}_t\left(\omega_{t+T}\right) + \theta\left[\mathbb{E}_t\left(\omega_{t+T}\right) - \mathbb{E}_{t-1}\left(\omega_{t+T}\right)\right]. \tag{3}$$

Furthermore, we have that $\mathbb{E}_t^\theta(\omega_{t+T}) = \mathbb{E}_t^\theta(\mathbb{E}_{t'}^\theta(\omega_{t+T}))$ for any $t < t' < t + T$.

Longer term forecasts can also be represented as a linear combination of past and present rational expectations. Furthermore, diagnostic expectations obey the law of iterated expectations with respect to the distorted expectations $\mathbb{E}_t^\theta$, so that forecast revisions are unpredictable from the vantage point of the decision maker.

However, forecast revisions are predictable using the true probability measure because errors in expectations correct on average in the future. Using (2), we find that

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

The average diagnostic forecast is rational. On average, diagnostic expectations revert to rational expectations because the diagnostic distortion is a linear function of news, and the average news is zero by definition. Even if expectations are inflated at $t-1$, they return to rationality on average at t. As we show in Section IV, this behavior allows us to account for the empirical findings in Section I.

III. A Model of Credit Cycles

We next introduce diagnostic expectations into a simple macroeconomic model and show that the psychology of representativeness generates excess volatility in expectations about credit spreads, overheating and overcooling of credit markets, and predictable reversals in credit spreads and economic activity that are consistent with the evidence in Section I, as well as with many other features of credit cycles.

[10] The Lucas critique holds that mechanical models of expectations cannot be used for policy evaluation because expectation formation in such models does not respond to changes in policy. Indeed, empirical estimates of mechanically extrapolative processes reveal parameter instability to policy change. Muth (1961) generalizes rational expectations to allow for systematic errors while preserving forward-looking behavior. His formula takes the linear form of equation (2): relative to rationality, expectations distort the effect of recent news. Muth's formulation naturally follows from the psychology of representativeness.

A. Production

A measure one of atomistic firms uses capital to produce output. Productivity at t depends on the state ω_t, but to a different extent for different firms. A firm is identified by its risk $\rho \in \mathbb{R}$. Firms with higher ρ are less likely to be productive in any state ω_t. If firm ρ enters period t with invested capital k, its current output is given by

$$y(k|\omega_t, \rho) = \begin{cases} k^\alpha & \text{if } \omega_t \geq \rho \\ 0 & \text{if } \omega_t < \rho \end{cases},$$

where $\alpha \in (0, 1)$. The firm produces only if it is sufficiently safe, $\rho < \omega_t$. The safest firms, for which $\rho = -\infty$, produce k^α in every state of the world. The higher is ρ, the better the state ω_t needs to be for the firm to pay off. At the same capital k, two firms produce the same output if they are both active, that is, if $\omega_t \geq \rho$ for both firms.

A firm's riskiness is common knowledge and is distributed across firms with density $f(\rho)$. Capital for production at $t+1$ must be installed at t, before ω_{t+1} is known. Capital fully depreciates after usage. At time t each firm ρ demands funds from a competitive financial market to finance its investment. The firm issues risky debt that promises a contractual interest rate $r_{t+1}(\rho)$. Debt is repaid only if the firm is productive: if at t the firm borrows $k_{t+1}(\rho)$ at the interest rate $r_{t+1}(\rho)$, next period it produces and repays $r_{t+1}(\rho)k_{t+1}(\rho)$ provided $\omega_{t+1} \geq \rho$, and defaults otherwise.

Because there are no agency problems and each firm's output has a binary outcome, the model does not distinguish between debt and equity issued by the firm. Both contracts are contingent on the same outcome and promise the same rate of return. For concreteness, we refer to the totality of capital invested as debt.

B. Households

A risk-neutral, infinitely lived representative household discounts the future by a factor $\beta < 1$. At each t, the household allocates its current income between nonnegative consumption and investment by maximizing its expectation of the utility function

$$\sum_{s=t}^{+\infty} \beta^{s-t} c_s.$$

The household purchases the claims issued by firms, which then pay out or default in the next period. The household's income consists of the payout of last period's debt, the profits of firms (which he owns), and a fixed endowment w that we assume to be large enough:[11]

[11] As we show below, this condition ensures that the equilibrium expected return is equal to β^{-1}.

ASSUMPTION 1: $w \geq (\alpha\beta)^{\frac{1}{1-\alpha}}$.

At each time s and state ω_s, the household's budget constraint is

$$c_s + \int_{-\infty}^{+\infty} k_{s+1}(\rho)\, f(\rho)\, d\rho = w + \int_{-\infty}^{+\infty} I(\rho, \omega_s)\left[r_s(\rho)\, k_s(\rho) + \pi_s(\rho)\right] f(\rho)\, d\rho,$$

where $c_s \geq 0$ is consumption, $k_{t+1}(\rho)$ is capital supplied to firm ρ, $I(\rho, \omega_s)$ is an indicator function equal to one when firm ρ repays, that is, when $\omega_s \geq \rho$, and $\pi_s(\rho)$ is the profit of firm ρ when active. The worse is the state of the economy (the lower is ω_s), the higher is the fraction of defaulting firms and thus the lower is the household's income.

The timeline of an investment cycle in the model is illustrated below.

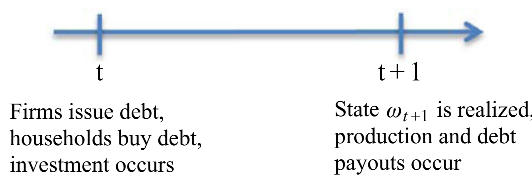

Investment by households and firms depends on the *perceived* probability with which each firm ρ repays its debt in the next period. At time t, the perceived probability with which firm ρ produces output and repays at time $t+1$ is given by

$$\mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right) = \int_{\rho}^{+\infty} h_t^{\theta}(x)\, dx = \frac{1}{\sigma\sqrt{2\pi}} \int_{\rho}^{+\infty} e^{-\frac{\left(x - \mathbb{E}_t^{\theta}(\omega_{t+1})\right)^2}{2\sigma^2}} dx.$$

The perceived probability of default is then $1 - \mu(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1}))$. A perfectly safe firm $\rho \to -\infty$ never defaults, since $\lim_{\rho \to -\infty} \mu(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})) = 1$.

When $\theta = 0$, expectations are rational and the probability of default is computed according to the true conditional distribution $h(\Omega_{t+1} = \omega | \Omega_t = \omega_t)$. When $\theta > 0$, the distortions of diagnostic expectations affect the perceived safety of different firms. In what follows, we refer to $\mu(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1}))$ as the "perceived creditworthiness" of firm ρ.

C. Capital Market Equilibrium and Credit Spreads

At time t, firm ρ demands capital $k_{t+1}(\rho)$ at the market contractual interest rate $r_{t+1}(\rho)$ so as to maximize its expected profit at $t+1$:

$$\max_{k_{t+1}(\rho)} \left(k_{t+1}(\rho)^{\alpha} - k_{t+1}(\rho) \cdot r_{t+1}(\rho)\right) \cdot \mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right).$$

The first order condition for the profit maximization problem is given by

$$k_{t+1}(\rho) = \left[\frac{\alpha}{r_{t+1}(\rho)}\right]^{\frac{1}{1-\alpha}}, \tag{4}$$

which is the usual downward sloping demand for capital.

Households are willing to supply any amount of capital to firm ρ provided that the interest rate $r_{t+1}(\rho)$ makes the household indifferent between consuming and saving:

$$r_{t+1}(\rho) \cdot \mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right) = \beta^{-1} \quad \Leftrightarrow \quad r_{t+1}(\rho) = \frac{1}{\beta \mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right)}. \tag{5}$$

In equilibrium, this condition must hold for all firms ρ. On the one hand, no arbitrage requires that all firms yield the same expected return. On the other hand, such expected return cannot be below β^{-1}. If this were the case, the household would not invest and the marginal product of capital would be infinite, leading to a contradiction. But the expected return of debt cannot be above β^{-1}. If this were the case, the household would invest the totality of its income. Under Assumption 1, however, this would imply that the marginal product of capital would fall below β^{-1}, again leading to a contradiction.

From equation (5), we can compute the spread obtained on the debt of risky firm ρ at time t as the difference between the equilibrium $r_{t+1}(\rho)$ and the safe rate β^{-1}:

$$S\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right) = \left(\frac{1}{\mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right)} - 1\right)\beta^{-1}. \tag{6}$$

Risky firms must compensate investors for bearing their default risk. The spread at t depends on the firm's riskiness ρ and on expectations of the aggregate economy. Greater optimism $\mathbb{E}_t^{\theta}(\omega_{t+1})$ improves perceived creditworthiness $\mu(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1}))$, lowering spreads. Greater riskiness ρ enhances spreads by reducing perceived creditworthiness.

Combining equations (6) and (4), we obtain

$$k_{t+1}(\rho) = \left[\frac{\alpha\beta}{1 + \beta S\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right)}\right]^{\frac{1}{1-\alpha}}, \tag{7}$$

which links expectations to credit spreads, investment, and output. In good times, households are optimistic, $\mathbb{E}_t^{\theta}(\omega_{t+1})$ is high, spreads are compressed, and firms issue debt and expand investment. When times turn sour, households become pessimistic, spreads rise, and firms cut debt issuance and investment. Equation (7) can be aggregated across different values of ρ to obtain aggregate investment at time t and output at $t+1$.

We can now generate testable implications of our model. Using equation (6), define the average spread at time t as

$$S_t = \int_{-\infty}^{+\infty} S\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right) f(\rho)\, d\rho. \tag{8}$$

Here, S_t is an inverse measure of optimism, which is strictly monotonically decreasing in expectations. Thus, when $\mathbb{E}_t^\theta(\omega_{t+1})$ is higher, average perceived creditworthiness is higher, and hence the average spread S_t charged on risky debt is lower.

We can substitute S_t for expectations $\mathbb{E}_t^\theta(\omega_{t+1})$ in equations (6) and (7). We then obtain the following result for the cross-section of firms.

PROPOSITION 2: *Lower optimism $\mathbb{E}_t^\theta(\omega_{t+1})$ and thus higher spread S_t at time t causes:*

(a) a disproportionate rise in the spread of riskier firms,

$$\frac{\partial^2 S\left(\rho, \mathbb{E}_t^\theta\left(\omega_{t+1}\right)\right)}{\partial S_t \partial \rho} > 0,$$

(b) a disproportionate decline in debt issuance and investment by riskier firms,

$$\frac{\partial}{\partial S_t} \frac{k_{t+1}\left(\rho_1\right)}{k_{t+1}\left(\rho_2\right)} < 0 \qquad \textit{for all } \rho_1 > \rho_2.$$

Because it is more sensitive to aggregate conditions, investment by riskier firms fluctuates more with expectations and displays more comovement with credit markets.

These predictions of the model are consistent with the evidence of Greenwood and Hanson (2013). They document that when the Baa-credit spread falls, bond issuance increases, with the effect particularly strong for firms characterized by higher expected default rates. As a consequence, the share of noninvestment grade debt over total debt (the "junk share") increases, as also documented by Lopez-Salido, Stein, and Zakrajsek (2017).[12]

This behavior of the junk share follows directly from property (b) above, which implies that the share of debt issued by firms riskier than an arbitrary threshold $\hat{\rho}$,

$$\frac{\int_{\hat{\rho}}^{+\infty} k_{t+1}\left(\rho\right) f\left(\rho\right) d\rho}{\int_{-\infty}^{+\infty} k_{t+1}\left(\rho\right) f\left(\rho\right) d\rho},$$

unambiguously increases as S_t drops and spreads become compressed (for any $\hat{\rho}$).

The qualitative effects described in Proposition 2 do not rely on diagnostic expectations and obtain even if households are fully rational. Diagnostic expectations have distinctive implications for the behavior of equilibrium credit spreads, as well as of their expectations by market participants, over time. We now turn to this analysis.

[12] Giroud and Mueller (2017) find that "high leverage firms" were worst hit during the 2008 financial crisis. Consistent with the above, such firms have most increased their leverage ratio, but are also less productive and more financially constrained than "low leverage firms." However, our stylized setting does not capture the fact that such firms also have higher leverage on average.

IV. Diagnostic Expectations and Equilibrium Credit Spreads

To study the equilibrium spread and expectations, we consider a linearized version of equation (6). A first-order expansion of equation (6) with respect to investors' expectations $\mathbb{E}_t^\theta(\omega_{t+1})$ around the long-run mean of zero yields

$$S\left(\rho, \mathbb{E}_t^\theta\left(\omega_{t+1}\right)\right) \approx \frac{1}{\beta} \cdot \left[\frac{1}{\mu\left(\rho, 0\right)} - 1\right] - \frac{\mu'\left(\rho, 0\right)}{\beta \mu(\rho, 0)^2} \cdot \mathbb{E}_t^\theta\left(\omega_{t+1}\right).$$

The spread drops as expectations improve (since $\mu'(\rho, 0) > 0$), but more so for riskier firms (the slope coefficient increases in ρ). Aggregating this equation across all firms ρ and denoting by $\sigma_0, \sigma_1 > 0$ the average intercept and slope, we find that the average spread at time t approximately satisfies

$$S_t = \sigma_0 - \sigma_1 \mathbb{E}_t^\theta\left(\omega_{t+1}\right). \tag{9}$$

Inserting into (9) the expression for $\mathbb{E}_t^\theta(\omega_{t+1})$ in equation (2), under the maintained assumption of AR(1) fundamentals $\omega_t = b\omega_{t-1} + \epsilon_t$, we establish the following result:

PROPOSITION 3: *The average credit spread S_t follows an ARMA(1,1) process given by*

$$S_t = (1-b)\,\sigma_0 + b \cdot S_{t-1} - (1+\theta)\, b\sigma_1\epsilon_t + \theta b^2\sigma_1\epsilon_{t-1}. \tag{10}$$

This is a key result. Under rational expectations ($\theta = 0$) the equilibrium spread, like fundamentals, follows an AR(1) process characterized by persistence parameter b. Starting from the long-run spread σ_0, a positive fundamental shock $\epsilon_t > 0$ causes expectations to improve and the spread to decline. Subsequently, the spread gradually reverts to σ_0. The reverse occurs after a negative piece of news $\epsilon_t < 0$: spreads go up on impact and then monotonically revert to σ_0.

Under diagnostic expectations, $\theta > 0$, credit spreads continue to have an autoregressive parameter b but now also contain a moving average component. The spread at time t now depends also on the shock experienced at $t-1$. If the news received in the previous period is good, $\epsilon_{t-1} > 0$, so that S_{t-1} is low, there is a discrete hike in the spread at time t. If the news received in the previous period is bad, $\epsilon_{t-1} < 0$, so that S_{t-1} is high, there is a discrete drop in the spread at time t.

These delayed corrections occur on average (controlling for mean-reversion in fundamentals) and correspond to the systematic correction of errors in diagnostic forecasts described in Section II. In fact, the overreaction $\theta\epsilon_{t-1}$ at $t-1$ reverses on average at t. Such reversal of expectations about fundamentals contaminates the spread, which exhibits the predictable nonfundamental reversal of equation (10). When at $t-1$ news is good, $\epsilon_{t-1} > 0$, optimism is excessive, and the spread S_{t-1} drops too far. Next period this excess optimism wanes on average, so that S_t is corrected upwards. The reverse occurs if at $t-1$ news is bad. In this sense, the making and un-making of expectational errors causes boom-bust cycles and mean-reversion in spreads.

We next show that the equilibrium behavior of credit spreads in (10) accounts for the findings on expectational errors of Section I. In Section IV.B, we show that the model also accounts for the evidence on the link between credit spreads and economic activity that is hard to explain under rational expectations.

A. Credit Spread Forecasts

In Section I, we show that forecasts of credit spreads exhibit predictable errors due to the extrapolative nature of expectations, and that these forecasts exhibit systematic reversals. To connect the model to this evidence, we now describe how agents with diagnostic expectations form forecasts of credit spreads.

Investors forecast future credit spreads using equation (9). As a consequence, the forecast made at t for the spread at $t+T$ is given by

$$\mathbb{E}_t^{\theta}\left(S_{t+T}\right) = \sigma_0 - \sigma_1 \mathbb{E}_t^{\theta}\left[\mathbb{E}_{t+T}^{\theta}\left(\omega_{t+T+1}\right)\right].$$

By exploiting Proposition 1 and Corollary 1, we obtain the following result.

LEMMA 1: *The T-periods-ahead diagnostic forecast of the spread is given by:*

$$\mathbb{E}_t^{\theta}(S_{t+T}) = \sigma_0\left(1 - b^T\right) + b^T S_t. \tag{11}$$

Diagnostic expectations project the current spread into the future via the persistence parameter b. The more persistent is the process for fundamentals, the greater is the influence of the current spread S_t on forecasts of future spreads.

Critically, unlike the equilibrium process for the spread in (10), the forecast process in (11) does not exhibit reversals. The intuition is simple: diagnostic forecasters fail to anticipate the systematic reversal in the equilibrium spread realized when their extrapolation of current news turns out to be incorrect.

This idea can help account for the evidence in Section I.

PROPOSITION 4: *If the equilibrium spread follows (10) and expectations follow (11), then:*

(a) the forecast error at $t+1$ is predictable in light of information available at t,

$$\mathbb{E}_t\left[S_{t+1} - \mathbb{E}_t^{\theta}\left(S_{t+1}\right)\right] = \theta b^2 \sigma_1 \epsilon_t, \tag{12}$$

(b) the revision of expectations about S_{t+T} occurring between t and $t+s$ is predictable in light of information available at time t,

$$\mathbb{E}_t\left[\mathbb{E}_{t+s}^{\theta}\left(S_{t+T}\right) - \mathbb{E}_t^{\theta}\left(S_{t+T}\right)\right] = \theta b^{T+1} \sigma_1 \epsilon_t. \tag{13}$$

Forecast errors and forecast revisions are predictable because agents neglect the reversal of their expectations errors. Thus, good news about fundamentals, $\epsilon_t > 0$, predicts that the realized spread next period is on average above the forecast (equation (12)) and that longer term forecasts of spreads will be revised upward in the future (equation (13)). The reverse pattern of predictability occurs after bad news $\epsilon_t < 0$. Equations (12) and (13) can therefore account for the evidence in Section I, whereby the spread S_t is negatively correlated with the forecast error $\mathbb{E}_t[S_{t+1} - \mathbb{E}_t^\theta(S_{t+1})]$ and with the future forecast revision $\mathbb{E}_t[\mathbb{E}_{t+s}^\theta(S_{t+T}) - \mathbb{E}_t^\theta(S_{t+T})]$.

We can draw a comparison between our model and "natural expectations" (Fuster, Laibson, and Mendel (2010)). In both models forecast errors are predictable because decision makers underestimate the possibility of reversals. The underlying mechanism, however, is very different. Under natural expectations, long-term reversals are assumed and errors in expectations arise because agents fit a simpler AR(1) model to the data. In our model, in contrast, both the process for spreads and forecast errors are endogenous to diagnostic expectations. Agents extrapolate current news too far into the future, which in turn endogenously generates unanticipated reversals in the equilibrium process for spreads.

To conclude, we use Proposition 3 and Lemma 1, together with the data on forecast errors analyzed in Section I, to provide a "back of the envelope" calibration of θ. This calculation should be taken with caution. Data on expectations of credit spreads are quite limited, available for only about 60 quarters. Moreover, our predictions also rely on a number of assumptions, including the assumption that fundamentals follow an AR(1) process. Still, it is useful to illustrate how data can be used to calibrate θ and thus to quantitatively discipline the model.[13]

LEMMA 2: *If fundamentals ω_t follow an AR(1) process with persistence b, the regression coefficient of spread forecast errors $S_{t+1} - \mathbb{E}_t^\theta(S_{t+1})$ on the current spread S_t is given by*

$$\gamma = -\frac{(1+\theta)\theta b}{(1+\theta)^2 + \frac{b^2}{1-b^2}}.$$

Regressing forecast errors on current levels of spreads yields an estimated coefficient that depends only on θ and the persistence b of fundamentals. In light of Lemma 1, persistence b can be estimated by regressing expectations $\mathbb{E}_t^\theta(S_{t+1})$ on current S_t. Table I presents results of both regressions, and provides estimates of γ (third column) and b (second column). From these estimates we find that $\theta = 0.91$.

This exercise suggests that forecast errors are sizable, in the sense that they are comparable to actual innovations in spreads—a point that is apparent from visual inspection of Figure 1. In the model, θ on the order

[13] A fuller quantification of θ based on field data requires developing tests that do not rely on exact knowledge of the data generating process. This lies beyond the scope of the current paper.

of one implies that investors effectively treat shocks ϵ_t to the state of the world as an AR(1) process with the same persistence as the underlying process. As a result, reversals in credit spreads $\theta b^2 \sigma_1 \epsilon_t$ are comparable in magnitude to the actual innovations $(1+\theta) b \sigma_1 \epsilon_{t+1}$. As we show next, this has significant implications for the predictability and excess volatility of returns.

B. Predictability of Returns, Volatility of Spreads, and Economic Activity

Our model can also account for the evidence on abnormal bond returns and excess volatility of credit spreads. To see this, define the "rational spread" S_t^r as one that would prevail at time t under rational expectations ($\theta = 0$). This is the compensation for default risk that rational investors demand. Proposition 3 then implies the following result.

COROLLARY 2: *Under diagnostic expectations, $\theta > 0$, the following properties hold:*

(a) investors earn predictably low (high) average returns after good (bad) news,

$$S_t - S_t^r = -\theta b \sigma_1 \epsilon_t,$$

(b) credit spreads exhibit excess conditional volatility,

$$Var\left[S_t \,|\omega_{t-1}\right] = (1+\theta)^2 Var\left[S_t^r \,|\omega_{t-1}\right].$$

Predictability of returns comes from errors in expectations. After good news $\epsilon_t > 0$, investors are too optimistic and demand too little compensation for default risk, $S_t < S_t^r$. The average realized return on bonds is thus below the riskless rate β^{-1}. After bad news $\epsilon_t < 0$, investors are too pessimistic and demand excessive compensation for default risk, $S_t > S_t^r$. The average realized return is thus above β^{-1}.[14]

Expectational errors also underlie the excess volatility of spreads. Equilibrium spreads vary too much relative to objective measures of default risk, which are captured by S_t^r, because spreads also reflect investor overreaction to recent news. Overreaction to good or bad news causes investors' risk perceptions to be too volatile, which introduces excess volatility into market prices.[15]

Greenwood and Hanson (2013) document the pattern of return predictability in Corollary 2. They find that high levels of the junk share predict anomalously low, and even negative, excess returns, and that this occurs precisely after good news, as measured by drops in expected default rates (point a). They

[14] Return predictability can also be gauged from equation (5). The average return at t is below the risk free rate when investors are too optimistic $\mu(\rho, \mathbb{E}_t^\theta(\omega_{t+1})) > \mu(\rho, \mathbb{E}_t(\omega_{t+1}))$, and above it otherwise.

[15] Excess volatility is due to the fact that beliefs do not depend only on the *level* of current fundamentals ω_t (as would be the case under rational expectations), but also on the magnitude ε_t of the recently observed news, which corresponds roughly speaking to the *change* in fundamentals.

consider conventional explanations for this finding, such as time-varying risk aversion and financial frictions, but conclude that the evidence (particularly the observed frequency of negative returns) is more consistent with the hypothesis that the junk share is a proxy for investor sentiment and extrapolation. Diagnostic expectations offer a psychological foundation for this account.

Several papers document that credit spreads appear too volatile relative to what could be explained by the volatility of default rates or fundamentals (Collin-Dufresne, Goldstein, and Martin (2001), Gilchrist and Zakrajšek (2012)). For instance, Collin-Dufresne, Goldstein, and Martin (2001) find that credit spreads display excess volatility relative to measures of fundamentals such as realized default rates, liquidity, or business conditions. They argue that this excess volatility—up to 75% of total realized volatility—can be explained by a common factor that captures aggregate shocks in credit supply and demand. In our model, investors' temporary overreaction to news about fundamentals offers a source for such shocks.

We conclude by illustrating a general implication of our model. Suppose, as we do, that the economy is driven by a single aggregate factor ω_t. With diagnostic expectations, regressing the average spread S_t on this aggregate factor would detect excess volatility.

LEMMA 3: *The R^2 of the regression of the average spread S_t on factor ω_t is given by*

$$R^2 = \frac{\left[1 + \theta\left(1 - b^2\right)\right]^2}{1 + 2\theta\left(1 - b^2\right) + \theta^2\left(1 - b^2\right)}.$$

Under rational expectations, $\theta = 0$, the spread is explained entirely by the level of fundamentals. Under diagnostic expectations, $\theta > 0$ (and provided there is persistence, $b > 0$), there is unexplained volatility in the spread, $R^2 < 1$. In fact, current fundamentals fail to capture reversals of past errors, which contaminate the spread dynamics. This feature does not arise under all belief distortions. For instance, if agents merely exaggerate the persistence of current fundamentals, changes in expectations are explained entirely by movements in ω_t, yielding an R^2 of one. In this sense, diagnostic expectations are a natural source of excess volatility.

The boom-bust cycles in credit spreads shape investment (see equation (7)) and lead in turn to overbuilding, underbuilding, and excess volatility in the real economy. Gennaioli, Ma, and Shleifer (2016) find that CFOs with more optimistic earnings expectations invest more. Greenwood and Hanson (2015) study investment cycles in the ship industry and find that, consistent with our model, returns to investing in dry bulk ships are predictable and tightly linked to boom-bust cycles in industry investment. High current ship earnings are associated with higher ship prices and higher industry investment, but predict low future returns on capital.

Our model also has implications for the link between credit markets and economic activity. Krishnamurthy and Muir (2015) and Lopez-Salido, Stein, and Zakrajsek (2017) document that a tightening of credit spreads at t induces output contraction in period $t+1$. Our model yields this pattern as a result of a drop in confidence. A reduction in optimism $\mathbb{E}_t^\theta(\omega_{t+1})$ raises the current spread. Tighter financial conditions lead current debt issuance and investment to decline, and in turn to a drop in aggregate output at $t+1$.

There is also growing evidence of systematic reversion in credit conditions and subsequent output declines. In particular, Lopez-Salido, Stein, and Zakrajsek (2017) show that low credit spreads at $t-1$ systematically predict higher credit spreads at t and then a drop in output at $t+1$. Lopez-Salido, Stein, and Zakrajsek (2017) suggest that a period of excessive investor optimism is followed by a period of cooling off, which they refer to as "unwinding of investor sentiment." This reversal contributes to a recession over and above the effect of changes in fundamentals. Diagnostic expectations can account for this unwinding of investor sentiment, thereby reconciling predictable reversals in market conditions with abnormal returns and excess volatility of credit spreads.[16] Such reversals are once again due to the autoregressive moving average (ARMA) process followed by the equilibrium spread in Proposition 3.

PROPOSITION 5: *Suppose that expectations are diagnostic, $\theta > 0$, and credit spreads are too low due to recent good news at t $-$ 1, that is, $\epsilon_{t-1} > 0$. Then:*

- *(a) controlling for fundamentals at t $-$ 1, credit spreads predictably rise at t, and*
- *(b) controlling for fundamentals at t $-$ 1, there is a predictable drop in aggregate investment at t and in aggregate production at t $+$ 1.*

Diagnostic expectations drive a cycle around fundamentals: overreaction to good news causes credit markets and the economy to overshoot at $t-1$. The subsequent reversal of such overreaction leads to a decrease in credit and economic activity that is more abrupt than could be accounted for by mean reversion in fundamentals (the result is fully symmetric for bad news, which leads to a bust followed by a boom). Indeed, investor psychology can itself be a cause of volatility in credit, investment, and business cycles, even in the absence of mean reversion in fundamentals, for example, if the process for aggregate productivity ω_t is a random walk ($b = 1$).

In sum, diagnostic expectations lead to short-term extrapolative behavior and systematic reversals. This is in line with a large set of recent empirical findings, including: (i) excess volatility of spreads relative to measures of fundamentals, (ii) excessive spread compression in good times and excessive spread

[16] In related settings, Jorda, Schularick, and Taylor (2013) document that strong growth of bank loans forecasts future financial crises and output drops. Baron and Xiong (2017) show that credit booms are followed by stock market declines. They document that in good times banks expand their loans, and this expansion predicts future negative returns on bank equity. The negative returns to equity may reflect the unwinding of initial investor optimism, or may be caused by abnormally low realized performance on the bank's credit decisions (as per Proposition 4). See also Fahlenbrach, Prilmeier, and Stulz (2016).

widening in bad times (and a similar pattern in the junk share), (iii) excessively volatile investment and output, (iv) good times predicting abnormally low returns, and (v) nonfundamental boom-bust cycles in credit spreads, driven by transient overreaction to news.

V. Safe Debt

So far, we have focused on the issuance and pricing of debt of different levels of risk, without introducing ingredients that make safe debt special. However, recent work highlights special features of safe debt, such as its attractiveness for investors demanding "money-like" securities (Stein (2014)). These features are important for thinking about financial crises. In these events, leverage takes center stage, giving rise to fire sales (when collateral is liquidated) or to a scarcity of safe assets (Gennaioli, Shleifer, and Vishny (2012, 2013)).

To illustrate the connection between safe debt, leverage, and diagnostic expectations, we assume that investors have a preferred habitat for safe debt. Specifically, investors require significantly lower returns for claims issued by firms that default with a probability lower than $\delta^* \in (0, 1)$. Gennaioli, Shleifer, and Vishny (2012) assume an extreme form of preferred habitat, modeled through infinite risk aversion.[17]

Preferred habitat creates a friction whereby financing of investment becomes easier when more firms are perceived to be safe. After good news, leverage expands across the economy but disproportionately so for those firms that are reclassified as safe. Conversely, contractions in safe debt cause a sudden tightening of financial constraints.

Investors value safe and risky debt separately.[18] When valuing safe debt, the investor trades off the cost of lending one dollar today to a risky firm ρ and receiving repayment $r_{t+1}(\rho)$ tomorrow in the case of no default using the usual discount factor β:

$$-1 + \beta r_{t+1}(\rho) \cdot \mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right). \tag{14}$$

When valuing risky debt, that is, a claim issued by a firm that defaults with believed probability $1 - \mu(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})) > \delta^*$, the investor instead applies the discount factor $\beta\psi$:

$$-1 + \beta\psi r_{t+1}(\rho) \cdot \mu\left(\rho, \mathbb{E}_t^{\theta}(\omega_{t+1})\right), \tag{15}$$

[17] Gennaioli, Shleifer, and Vishny (2012) also consider the case in which only some investors have a preference for safe debt. Preference heterogeneity can help understand phenomena such as trading of financial assets. For simplicity, we continue to assume that there is a representative investor, as in the previous sections.

[18] The fact that an investor values securities one by one, rather than valuing an overall portfolio, reflects the idea of preferred habitat: some investors specialize in holding a certain asset type. Without this assumption, even risk-averse investors with a preference for low default risk would be willing to hold some risky assets as long as the overall default probability of their portfolio is below δ^*.

where $\psi \leq 1$ captures the idea that the investor requires a higher return to hold risky debt (which can also be thought of as equity). A preference for safe debt arises when $\psi < 1$ and $\delta^* < 1$. The model boils down to infinite risk aversion for $\psi = \delta^* = 0$.

We can now describe the equilibrium.

PROPOSITION 6: *At time t define a threshold ρ_t^θ as the solution to*

$$\mu\left(\rho_t^\theta, \mathbb{E}_t^\theta\left(\omega_{t+1}\right)\right) = 1 - \delta^*,$$

so that ρ_t^θ increases in optimism about the future $\mathbb{E}_t^\theta(\omega_{t+1})$. We then have:

(i) firms with $\rho \leq \rho_t^\theta$ issue safe debt, on which they promise the equilibrium interest rate

$$r_{t+1}(\rho) = [\beta\mu(\rho, \mathbb{E}_t^\theta(\omega_{t+1}))]^{-1},$$

(ii) firms with $\rho > \rho_t^\theta$ issue risky debt, on which they promise the equilibrium interest rate

$$r_{t+1}(\rho) = [\beta\psi\mu(\rho, \mathbb{E}_t^\theta(\omega_{t+1}))]^{-1}.$$

Investors' preference for safe debt creates a discontinuity in required returns. Safe firms can borrow at an expected return of β^{-1}, which is lower than the expected return $(\beta\psi)^{-1}$ that risky firms must pay. The threshold ρ_t^θ that separates safe from risky firms is time-varying. As more firms are categorized as safe in good times, the required return for marginally risky firms declines discontinuously.

In good times, the issuance of safe debt expands along both the intensive and the extensive margins. On the intensive margin, already safe firms issue too much safe debt, relative to the rational expectations benchmark. On the extensive margin, too many risky firms become safe, and start issuing safe debt (this occurs when ρ_t^θ is too high relative to the level $\rho_t^{\theta=0}$ that would prevail under rational expectations). In this sense, at the end of expansions borrowers (particularly marginally risky ones) are excessively burdened with liabilities perceived to be safe.

This mechanism implies that after periods of good news the debt structure is fragile, as in Gennaioli, Shleifer, and Vishny (2012). Diagnostic expectations imply that crises occur when good news stops coming, so that excess optimism reverts. To see this, consider a thought experiment outside the current model. Suppose that, at an interim period between time t and time $t+1$, investors reassess their expectations of future economic conditions ω_{t+1}, which revert to rational expectations (e.g., because no news was received, in line with Proposition 1). We then have the following result.

PROPOSITION 7: *After good news, $\epsilon_t > 0$, the interim valuation of assets moves as follows:*

(a) *The valuation of safe and risky debt, that is, of firms with* $\rho \notin (\rho_t^{\theta=0}, \rho_t^{\theta})$*, drops by*

$$\frac{\mu\left(\rho, \mathbb{E}_t^{\theta}\left(\omega_{t+1}\right)\right)}{\mu\left(\rho, \mathbb{E}_t^{\theta=0}\left(\omega_{t+1}\right)\right)} > 1.$$

(b) *The valuation of debt perceived to be safe, that is, of firms with* $\rho \in (\rho_t^{\theta=0}, \rho_t^{\theta})$*, drops by*

$$\frac{\mu\left(\rho, \mathbb{E}_t^{\theta}\left(\omega_{t+1}\right)\right)}{\psi\mu\left(\rho, \mathbb{E}_t^{\theta=0}\left(\omega_{t+1}\right)\right)} \gg 1.$$

Because in the interim period there is no new information, under rational expectations the valuation of different securities would not change. Instead, under diagnostic expectations the initial excess optimism wanes and the probability of repayment of all firms is reassessed downward, leading to negative realized returns.[19]

The key result in Proposition 7 is that realized returns are particularly low for the marginally risky firms $\rho \in (\rho_t^{\theta=0}, \rho_t^{\theta})$. These firms are considered safe under the initial excess optimism, but not when investors appreciate their risk. Investors then demand a higher expected return to hold this debt, causing its price to exhibit a sharp decline.

As in Gennaioli, Shleifer, and Vishny (2012), financial fragility arises from the combination of neglect of risk and preferred habitat. In the absence of preferred habitat (e.g., if $\psi = 1$), diagnostic expectations boost risk taking in good times as well as overvaluation and predictability in bond returns, but in a continuous way. A preferred habitat for safe debt, in contrast, enhances volatility and concentrates it precisely on the safe debt segment. The overissuance of safe debt causes neglected risks to be sharply underpriced, which exposes the economy to bad shocks when neglected risks resurface and safe debt is reclassified as risky. Even small downward revisions of expected fundamentals can trigger a large price decline when it causes a reassessment of safety. In line with the discussion in Section II, unintended exposure to neglected risks is larger when economic volatility σ^2 is lower.

The price effects on debt described above can percolate through the economy in many ways. And while the price mechanism is symmetric with respect to good and bad news, its impact on the economy may be highly skewed. For example, if safe debt is an asset on the balance sheets of intermediaries, these intermediaries might become insolvent, or experience runs, when the price of debt falls (Baron and Xiong (2017), Fahlenbrach, Prilmeier, and Stulz (2016)). This can have large adverse effects on the economy through the bank lending channel (Chodorow-Reich (2014)). Leveraged investors holding such "safe" debt

[19] The result carries through under the alternative assumption that information also arrives in the interim period. However, the analysis would be more complicated because (i) ex post valuation would also change under rationality, since new information has arrived, and (ii) ex ante valuation would take into account the probability that firms become safer or riskier than δ^* before repayment.

might have to sell it at fire-sale prices, which can adversely affect balance sheets of other intermediaries or firms. Such compounding of adverse effects does not have a counterpart for positive news, creating a significant real asymmetry in response to movements in debt prices.

VI. Conclusion

In this paper, we present a new approach to modeling beliefs in economic models—diagnostic expectations—that is based on Kahneman and Tversky's (1972) representativeness heuristic. Our model of expectations is portable in Rabin's (2013) sense, meaning that the same framework accounts for many experimental findings, the phenomenon of stereotyping, but also critical features of beliefs in financial markets such as extrapolation, overreaction, and neglect of risk. Diagnostic expectations are also forward looking, which means that they are invulnerable to the Lucas critique of mechanical backward looking models of beliefs. We apply diagnostic expectations to a straightforward macroeconomic model of investment, and find that it can account for several empirical findings regarding credit cycles without resorting to financial frictions.

Initial submission: July 6, 2016; Accepted: February 22, 2017
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Arnold, Marc, Dustin Schuette, and Alexander Wagner, 2015, Pay attention or pay extra: Evidence on the compensation of investors for the implicit credit risk of structured products, Research paper no. 14–24, Swiss Finance Institute.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2016, Extrapolation and bubbles, *Journal of Financial Economics* (forthcoming).

Barberis, Nicholas, and Andrei Shleifer, 2003, Style investing, *Journal of Financial Economics* 68, 161–199.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk, *Quarterly Journal of Economics* 132, 765–809.

Bernanke, Ben, 1990, On the predictive power of interest rates and interest rate spreads, *New England Economic Review*, November/December, 51–68.

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

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

Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar, 2016, Inattentive Expectations and Stock Market Anomalies, Working Paper MIT.

Cagan, Phillip, 1956, The monetary dynamics of hyperinflation, in Milton Friedman, ed.: *Studies in the Quantity Theory of Money* (University of Chicago Press, Chicago, IL).

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

Chodorow-Reich, Gabriel, 2014, The employment effects of credit market disruptions: Firm-level evidence from the 2008–09 financial crisis, *Quarterly Journal of Economics* 129, 1–59.

Coibion, Olivier, and Yuriy Gorodnichenko, 2012, What can survey forecasts tell us about informational rigidities?, *Journal of Political Economy* 120, 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.

Collin-Dufresne, Pierre, Robert Goldstein, and Spencer Martin, 2001, The determinants of credit spread changes, *Journal of Finance* 56, 2177–2207.

Coval, Joshua, Kevin Pan, and Erik Stafford, 2014, How markets learn, Working paper, Harvard Business School.

Cutler, David, James Poterba, and Lawrence Summers, 1990, Speculative dynamics and the role of feedback traders, *American Economic Review Papers & Proceedings* 80, 62–68.

DeLong, Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45, 379–395.

Eggertsson, Gauti, and Paul Krugman, 2012, Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach, *Quarterly Journal of Economics* 127, 1469–1513.

Fahlenbrach, Rudiger, Robert Prilmeier, and René Stulz, 2016, Why does fast loan growth predict poor performance for banks? NBER Working paper 22089.

Friedman, Benjamin, and Kenneth Kuttner, 1992, Money, income, prices, and interest rates, *American Economic Review* 82, 472–492.

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

Gabaix, Xavier, 2014, A sparsity-based model of bounded rationality, *Quarterly Journal of Economics* 129, 1661–1710.

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, *Quarterly Journal of Economics* 125, 1399–1433.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2012, Neglected risks, financial innovation, and financial fragility, *Journal of Financial Economics* 104, 452–468.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2013, A model of shadow banking, *Journal of Finance* 68, 1331–1363.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2015, Neglected risks: The psychology of financial crises, *American Economic Review, Papers and Proceedings* 105, 310–314.

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

Giroud, Xavier, and Holger Mueller, 2017, Firm leverage, consumer demand, and employment losses during the Great Recession, *Quarterly Journal of Economics* 132, 271–316.

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

Greenwood, Robin, and Samuel G. Hanson, 2015, Waves in ship prices and investment, *Quarterly Journal of Economics* 130, 55–109.

Greenwood, Robin, Samuel G. Hanson, and Lawrence Jin, 2017, A model of credit market sentiment, Working paper no. 2016–02, Harvard Business School Project on Behavioral Finance and Financial Stability.

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

Hirshleifer, David, Jun Li, and Jianfeng Yu, 2015, Asset pricing in production economies with extrapolative expectations, *Journal of Monetary Economics* 76, 87–106.

Hong, Harrison, and Jeremy Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jin, Lawrence, 2015, A speculative asset pricing model of financial instability, Working paper, Yale School of Management.

Jorda, Oscar, Moritz Schularick, and Alan Taylor, 2013, When credit bites back: Leverage, business cycles, and crises, *Journal of Money, Credit and Banking* 45, 3–28.

Kahneman, Daniel, 2011, *Thinking: Fast and Slow* (Farrar, Straus, and Giroux, New York, NY).

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

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Krishnamurthy, Arvind, and Tyler Muir, 2015, Credit spreads and the severity of financial crises, Working paper, Stanford University.

Lopez-Salido, David, Jeremy Stein, and Egon Zakrajsek, 2017, Credit-market sentiment and the business cycle, *Quarterly Journal of Economics* 132, 1373–1426.

Mian, Atif, Amir Sufi, and Emil Verner, 2017, Household debt and business cycle worldwide, *Quarterly Journal of Economics* 132, 1755–1817.

Minsky, Hyman, 1977, The financial instability hypothesis: An interpretation of Keynes and an alternative to "standard" theory, *Nebraska Journal of Economics and Business* 16, 5–16.

Muth, John, 1961, Rational expectations and the theory of price movements, *Econometrica* 29, 315–335.

Newey, Whitney, and Kenneth West, 1994, Automatic lag selection in covariance matrix estimation, *Review of Economic Studies* 61, 631–653.

Rabin, Matthew, 2013, Incorporating limited rationality into economics, *Journal of Economic Literature* 51, 528–543.

Rabin, Matthew, and Dimitri Vayanos, 2010, The gambler's and hot hand fallacies: Theory and applications, *Review of Economic Studies* 77, 730–778.

Rognlie, Matthew, Andrei Shleifer, and Alp Simsek, 2017, Investment hangover and the Great Recession, *American Economic Journal: Macroeconomics* (forthcoming).

Schularick, Moritz, and Alan Taylor, 2012, Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008, *American Economic Review* 102, 1029–1061.

Sims, Christopher, 2003, Implications of rational inattention, *Journal of Monetary Economics* 50, 665–690.

Stein, Jeremy, 2014, Incorporating financial stability considerations into a monetary policy framework, Speech at the International Research Forum on Monetary Policy, Federal Reserve Board of Governors, Washington, D.C., March 21.

Stock, James, and Mark Watson, 2003, Forecasting output and inflation: The role of asset prices, *Journal of Economic Literature* 41, 788–829.

Tversky, Amos, and Daniel Kahneman, 1983, Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment, *Psychological Review* 90, 293–315.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.