

The Cushioning Benefits of Biased Beliefs

Lawrence J. Jin, Matthew S. Shum, Mali Zhang*

July 4, 2018

ABSTRACT

We present a model of dynamic investment and production in which producers may have biased beliefs in which they overextrapolate recent demand conditions into the future. This bias leads producers' beliefs to exhibit *insufficient mean reversion*, as these producers underestimate the degree of mean reversion in the demand process. In a volatile industry, while biased beliefs lead firms to make sub-optimal investment decisions in the short-run, they can be beneficial in the long-run by counteracting the general trend in the industry, “cushioning” the industry against prolonged downturns and aiding faster recovery. As an empirical case study, we consider oil exploration in Alaska. We present evidence that firms in this industry were subject to extrapolation bias, leading to drilling of lower-profit wells after recent price increases. Calibration of our model to Alaska oil exploration shows that the cushioning effect can be large: in a typical episode of oil price decline arising from a sequence of adverse demand shocks, the cushioning effect reduces the decline of the oil price by 8.2% and accelerates the price recovery by 27%. This showcases the potential positive implications that biased beliefs can have on industry dynamics.

*We thank Nicholas Barberis, Cary Frydman, Robin Greenwood, Ryan Kellogg, Charles Nathanson, Tom Ruchti, Andrei Shleifer, and participants at the Caltech Brown Bag for helpful comments. All three authors' affiliation is the Division of the Humanities and Social Sciences at the California Institute of Technology. Their email addresses are lawrence.jin@caltech.edu, mshum@caltech.edu, and mzzhang@caltech.edu, respectively.

Biased belief formation is a bedrock underlying many models in behavioral finance and behavioral economics. In financial settings, many real-world investors, both individual and institutional investors, exhibit extrapolative expectations: they believe that prices will remain high in the future when recent prices have been high. Accordingly, these beliefs can be conducive to price bubbles, and sometimes tend to lengthen the duration of bubbles (Bagehot (1873), Kindleberger (1978), Barberis, Greenwood, Jin, and Shleifer (2017)). Conversely, extrapolative expectations can amplify downward price movements and sometimes lead to slow recovery in an industry or economy (Greenwood, Hanson, and Jin (2016), Jin (2015)). By and large, biased beliefs have negative implications for market dynamics in these settings.

A large literature has focused on consumers having biased beliefs and how their beliefs affect trading in asset markets. In reality, however, biased beliefs arguably also play a role in many “real” investment decisions. Indeed, some recent research explores how biased expectations can impact such decisions. For instance, Gennaioli, Ma, and Shleifer (2015) document that corporate investment plans and actual investments are explained by CFOs’ incorrect expectations; Greenwood and Hanson (2015) study how biased beliefs generate return predictability in the global ship building industry.

In this paper, we study a model of biased beliefs populated by producers who make real investment decisions based on these beliefs. Specifically, the biased beliefs take the form of backward-looking extrapolation: producers’ expectations of future consumer demand is formed as a weighted average of these consumers’ past demands.¹ These beliefs exhibit *insufficient mean reversion*: producers mistakenly assume that the long-run mean of the demand process is changing and estimate it using recent realizations of demand, thus underestimating the degree of mean reversion in the process. These biased beliefs affect producers’ investment behavior, leading to sub-optimal decisions. Obviously, sub-optimal decisions impose welfare losses on the economy, and much of the existing behavioral literature focuses on these negative effects arising from agents’ behavioral biases.

In this paper, however, we go one step further. When firms’ investment decisions occur within a dynamic market equilibrium, these “mistakes” can actually translate into long-run gains in the market. Specifically, we show that over time, biased beliefs generate some unexpected effects that counteract the general trend of an industry or economy, “cushioning” the industry or economy against prolonged downturns and therefore leading to faster recovery. During industry upturns, these cushioning effects can shorten the duration of bubbles. Unlike many papers in the existing literature, our focus is on the *positive* implications that biased beliefs may play in market dynamics.

To illustrate these cushioning benefits, consider the oil exploration industry, an industry with pronounced boom-and-bust cycles and volatile prices in which extrapolation can have big effects. In this market, the producers are large oil companies who make important decisions with long-run impact on oil exploration and production. When oil prices spiral downward (as occurred recently

¹Earlier works of Barberis and Shleifer (2003), Barberis, Greenwood, Jin, and Shleifer (2015), and Hirshleifer, Li, and Yu (2015) propose extrapolative beliefs about stock market returns and GDP growth.

in the world oil market), oil companies extrapolate low prices continuing into the future, and therefore cut back on new exploration. Indeed, extrapolative beliefs cause firms to cut exploration *more* than they would in the non-extrapolative benchmark, resulting in large welfare losses in the interim. However, over time, this excessive reduction in oil production will put *upward* pressure on prices, thus reversing the downward trend in prices and aiding the oil industry out of its doldrums. Moreover, this recovery will happen faster when firms have extrapolative beliefs. Conversely, in periods of rising prices, extrapolative producers overinvest in oil exploration, which puts downward pressure on the rising prices. We show that these cushioning effects constitute a generic feature of real investment models with producers having backward-looking extrapolative beliefs.

Our model builds on standard aggregate investment models of [Abel \(1981\)](#) and [Abel and Eberly \(1994\)](#). The price of industry output is positively related to consumer demand and negatively related to total investment from producers. Over time, consumer demand follows a mean-reverting process with a constant long-run mean. Without knowing this long-run mean, however, producers extrapolate past realizations of consumer demands in forming expectations of future demand. Based on these beliefs, producers make investment decisions. For comparison, we also examine a benchmark model in which producers know the long-run mean of consumer demand. To analyze the cushioning effects, we analyze the impulse response of investment, total supply, and product prices with respect to shocks to consumer demands. Our analysis leads to two observations. First, with extrapolative beliefs from producers, a negative demand shock gives rise to persistent underinvestment in subsequent periods, causing total supply to decrease at a faster rate relative to the benchmark model. This rundown in supply lends support to prices, thus “cushioning” the negative impact of the demand shock on the product price. Second, due to the persistence of the cushioning effects, the product price can sometimes even start rising in the midst of a sequence of negative demand shocks.

As an empirical case study, we consider the behavior and experience of oil producers operating in the North Slope of Alaska, one of the most active oil exploration sites in North America. Since oil exploration is not a liquid asset market with ample trading and resale opportunities, the return regressions or survey evidence used in the existing literature to detect extrapolation are not available in this context. Given these challenges, we present several pieces of evidence from Alaska exploration which are consistent with the presence of biased beliefs on the part of producers. First, we find that the number of new wells drilled are positively correlated with past levels of oil prices, with more significant correlation with prices from six to twelve months prior. Second, we find that the five-year production and five-year revenue from oil production for newly drilled wells are both *negatively* correlated with past levels of oil prices. This finding that drilling projects initiated following high prices yield systematically inferior outcomes support our interpretation of these projects as resulting from biased beliefs.

As a further test for firms’ biased beliefs, we exploit the availability of data on both oil companies’ planned and actual investments to establish that the number of “scrapped” wells — that

is, the difference between the number of wells actually drilled and the number of wells planned — is negatively correlated with both past levels of oil prices and the average changes in prices during the subsequent two-year period within which oil companies are approved to carry out the planned drilling. This finding suggests that oil companies over-extrapolate past oil prices when planning for well drilling but subsequently change their mind and forgo these opportunities if oil prices decline during the time after the well was approved but before drilling has commenced.²

In addition, we document some path-dependent features of belief formation. When recent oil prices have risen at an increasing rate, oil companies extrapolate to a lesser extent; if the oil price six months ago is 50 dollars per barrel, then all else being equal, an increasing rise of the oil price between six months ago and the time of drilling will result in a 12% reduction of the number of new wells drilled.

Motivated by this empirical evidence of oil companies' suboptimal exploration decisions in response to recent prices, we quantify the magnitude of the cushioning benefit by calibrating our model using parameters appropriate to this industry. In one example, a sequence of adverse demand shocks leads to a price decline which is 8.2% smaller in the extrapolative compared to the benchmark scenario, which quickens the recovery by four months. Another calibration example suggests that the industry downturn during the 2008 financial crisis would have been lengthened by four months if oil industry firms did *not* have biased extrapolative beliefs. These examples show how extrapolative beliefs cushion, or soften, the extremes during the periodic big downturns which punctuate the oil exploration industry. At the same time, the overall welfare calculus is ambiguous; while this cushioning effect shortens the downturn, the biased beliefs lead firms to severely underinvest in drilling activity during the downturn relative to the non-extrapolative benchmark, which involves large costs in lost jobs, underutilized equipment, and so on. Biased beliefs lead to *shorter, albeit direr*, downturns, and the overall welfare effect involves a tradeoff between the short-run costs and the long-run gains.

Our paper adds to both theoretical and empirical research that aims to understand the implications of biased beliefs on asset price movements, consumption and portfolio choices, investment decisions, and individual behavior. On the empirical side, recent papers by [Vissing-Jorgensen \(2004\)](#), [Bacchetta, Mertens, and van Wincoop \(2009\)](#), [Amromin and Sharpe \(2013\)](#), [Greenwood and Shleifer \(2014\)](#), [Kojien, Schmeling, and Vrugt \(2015\)](#), and [Kuchler and Zafar \(2016\)](#) present survey evidence that real-world investors exhibit extrapolative expectations and they behave according to these beliefs. [Cassella and Gulen \(2015\)](#) show that the extent to which investors extrapolate past returns of the stock market is highly correlated with the degree of predictability of future market returns. And [Gennaioli et al. \(2015\)](#) find that CFO expectations about future earnings growth are extrapolative and predictive of planned and actual investments. On the theoretical side, [Fuster, Hebert, and Laibson \(2011\)](#), [Choi and Mertens \(2013\)](#), [Hirshleifer et al. \(2015\)](#), [Barberis et al. \(2015\)](#), [Jin and Sui \(2017\)](#) show that extrapolative expectations can generate stock market move-

²Similarly, [Conlin, O'Donoghue, and Vogelsang \(2007\)](#) used data on purchased and subsequently returned clothing to identify projection bias.

ments that are consistent with the data. [Alti and Tetlock \(2014\)](#) show that overextrapolation and overconfidence affect investment decisions. [Barberis et al. \(2017\)](#) use extrapolation to explain asset bubbles and trading volume. [Gennaioli, Shleifer, and Vishny \(2012\)](#) and [Jin \(2015\)](#) connect biased beliefs that arise from risk neglect and availability heuristic with market leverage and financial crashes. [Glaeser and Nathanson \(2015\)](#) tie extrapolation to housing dynamics. And [Bordalo, Gennaioli, and Shleifer \(2017\)](#) and [Greenwood et al. \(2016\)](#) use extrapolative expectations to make sense of facts about credit cycles.

Our paper makes three contributions to this line of research. First, the paper embeds extrapolative expectations into the supply side of a model with real investments and unearths the potential cushioning benefits from these biased beliefs. It is worth pointing out that most research so far has focused on the negative effects and costs associated with biased beliefs; only [Dong, Hirshleifer, and Teoh \(2017\)](#) discuss the potential benefits. Our study adds to the latter strand of literature and counsels greater caution when assessing the overall effect of biased beliefs. Second, we present empirical evidence from the oil exploration industry that supports our model assumptions and predictions. Thus our paper joins a small but growing literature exploring and quantifying the impact of biased beliefs in a specific industrial (i.e., non-financial) setting.³ Our data allow us to look at the planned and actual investments separately, and studying the difference between the two supports the hypothesis that producers on the supply side exhibit biased beliefs. Finally, we use our data to calibrate model parameters; this allows us to further quantify the cushioning effects highlighted in the paper.

Our paper is related to the works of [Greenwood and Hanson \(2015\)](#) and [Bordalo et al. \(2017\)](#); these works also connect biased beliefs with investment decisions. Different from their studies which analyze the asset price implications of biased beliefs, we focus on the cushioning benefits of these beliefs. Our paper is also related to the work of [Glaeser, Gyourko, and Saiz \(2008\)](#). Their work shows that the elasticity of housing supply can affect the magnitude and duration of housing bubbles, and biased beliefs in their framework come from the demand side. Instead, our study highlights the importance of biased beliefs from the supply side.

The paper proceeds as follows. In Section [I](#), we lay out the model and characterize its solution. We then use impulse responses of the model to illustrate the cushioning benefits. Section [II](#) uses the Alaska exploration data to provide evidence that supports the extrapolative bias amongst oil and gas companies. In Section [III](#), we calibrate model parameters in accordance with the data and further analyze the implications of the model using these parameter values. Section [IV](#) concludes. All technical details are in the Appendix.

³The other paper is [Greenwood and Hanson \(2015\)](#). [Kellogg \(2014\)](#) estimates a structural model of individual oil companies' oil drilling decisions in Texas and also estimates producers' belief process for future oil prices.

I. The Model

In this section, we first develop a simple aggregate investment model with incorrect beliefs from producers. Then we examine the model implications through impulse response analysis. For congruence with the empirical case study which follows, we will describe the model using terminology from the oil industry. The firms are oil producers who make decisions about the number of wells to drill each period, and obtain revenue from selling the oil extracted from the wells.

Assume that the demand relationship between crude oil prices per barrel H_t and the total number of active wells Q_t is

$$H_t = A_t - BQ_t. \quad (1)$$

Here A_t represents a demand factor; this captures outside influences on prices which are exogenous to the firms. Since the oil market is global, these influences can include supply disturbances in other oil-producing areas of the world (such as Texas, Canada, the Middle East, etc.) which will also impact the price that Alaskan producers receive for their oil. Such disturbances evolve randomly and with some serial dependence, so we model the law of motion for A_t as

$$A_{t+1} = \bar{A} + \rho_0(A_t - \bar{A}) + \varepsilon_{t+1}, \quad (2)$$

with $\rho_0 \in [0, 1)$ and $\text{Var}[\varepsilon_{t+1}] = \sigma_\varepsilon^2$. Q_t , the aggregate number of wells, is an investment decision made by a continuum of risk-neutral firms. At each point in time, each firm chooses its level of investment i_t^G . The relation between the firm's wellcount q_t and its time- t investment is

$$q_{t+1} = (1 - \delta)q_t + p \cdot i_t^G = q_t + p \cdot i_t, \quad (3)$$

where $0 < \delta < 1$ is the depreciation rate, p is the probability of success when producing the industry output, and $i_t = i_t^G - \delta q_t/p$.⁴ Effectively, i_t is the choice variable. At the aggregate level, the total number of wells evolves as

$$Q_{t+1} = (1 - \delta)Q_t + p \cdot I_t^G = Q_t + p \cdot I_t. \quad (4)$$

This law of motion for Q_t implies that there is a one period time-lag in investment I_t before it affects the number of wells Q_{t+1} , and generates cash flow for the firms. In mapping this model to the oil exploration industry, we use a period of a month, which is reasonable given the lag between drilling and well production falls between a few days and a couple of months for most wells in our data.⁵

⁴Under risk neutrality and the assumption that success or failure of production is independent across firms, (3) is equivalent to leaving the incremental investment stochastic and then taking expectations when deriving the Bellman value function.

⁵In other industries, such as the housing market, it may take much longer (perhaps years rather than months) for new planned housing to be completed, and we conjecture that with such long delays, biased beliefs may actually

Similar to [Greenwood and Hanson \(2015\)](#), we assume that the representative firm earns a net profit of

$$\Pi_t = M \cdot (A_t - BQ_t) - C - \delta P_r \quad (5)$$

on each active well, where M is the average number of barrels obtained from each well, C is the operating cost of a well, and P_r is the replacement cost of a well. For an individual firm, given its current wellcount q_t and its current investment i_t (new drilling), the firm’s time- t total profit is

$$V_t = q_t \Pi_t - P_r i_t - k \cdot \frac{i_t^2}{2}, \quad (6)$$

where $k \cdot \frac{i_t^2}{2}$ represents the adjustment costs.

Firms’ biased beliefs and insufficient mean reversion. A crucial component of our model lies in the specification of firms’ expectational errors. Specifically, we assume that, from firms’ perspective, the evolution of the demand factor A_t is

$$A_{t+1} = \bar{A}_t^\alpha + \rho_f \cdot (A_t - \bar{A}_t^\alpha) + \varepsilon_{f,t+1}, \quad (7)$$

where

$$\bar{A}_t^\alpha = \alpha \cdot \bar{A}_t + (1 - \alpha) \bar{A}, \quad \bar{A}_t = (1 - \rho_A) \bar{A}_{t-1} + \rho_A A_t, \quad (8)$$

the subscript “ f ” stands for “firm”, and $1 > \rho_f \geq \rho_0 \geq 0$.⁶

Comparing Equation (2) to Equations (7) and (8) shows that such a model of beliefs exhibits *insufficient mean reversion*; \bar{A} is the true long-run mean of the demand process, which is constant over time, but firms mistakenly assume it to be time-varying (denoted \bar{A}_t^α in Equation (8)), and estimate it using recent realizations of the process. Thus, in each period, firms’ beliefs about the long-run mean of the process adjust in the direction of recent realizations: when demand has been slack—that is, when A_t has been low—firms tend to believe that the long-run mean \bar{A}_t^α has also fallen, leading to a smaller perceived degree of mean reversion measured by the difference between \bar{A}_t^α and A_t ; the opposite occurs after periods when demand has been high.

[Place Figure 1 about here]

Figure 1 contains an illustration of these beliefs. The demand factor process, A_t , is plotted in the solid line, and there is first an upturn followed by a downturn. However, the true long-run mean, \bar{A} , is constant over time and equal to zero in this example, as plotted in the dotted line. In contrast, firms’ beliefs, following Equations (7) and (8), are characterized by a time-varying long-run mean, \bar{A}_t^α , which is plotted in the dot-dash line. Clearly, firms’ beliefs about the long-run tendency of the

exacerbate rather than cushion the economy against downturns.

⁶Note that when $\alpha = 0$, \bar{A}_t^α equals \bar{A} . In this case, the model reduces to the model of [Greenwood and Hanson \(2015\)](#). In comparison to [Greenwood and Hanson \(2015\)](#), our way of modelling extrapolative expectations allows us to further make sense of some path-dependent features of belief formation; we discuss this both later in this section and in Section II.

process exhibit insufficient mean reversion: following an upturn in A_t , firms' perceived long-run means also track higher, and as a result, the perceived degree of mean reversion in A_t , measured by $A_t - \bar{A}_t^\alpha$, becomes smaller, suggesting "irrational exuberance", while the opposite occurs following the downturn in A_t , suggesting "irrational pessimism".

In addition, conditional on their estimated long-run mean \bar{A}_t^α , we can also allow firms to perceive less mean reversion of A_t relative to the true process by having $\rho_f \geq \rho_0$. Rearranging (7)

$$A_{t+1} - A_t = (1 - \rho_f) (\bar{A}_t^\alpha - A_t) + \varepsilon_{f,t+1} \quad (9)$$

illustrates the path-dependent feature of firms' belief formation. With a sequence of steady increase in the demand factor, \bar{A}_t^α rises above the true long-run mean of \bar{A} . In this case, a high \bar{A}_t^α and a high ρ_f both make the perceived evolution of A_{t+1} less mean-reverting than the true data generating process. If, on the other hand, A_t rises at an increasing rate, then \bar{A}_t^α increases to a smaller degree compared to A_t , making firms perceive that A_t will mean-revert back to a low level more quickly in the future.

Equations (7) and (8) are also related to the work of [Barsky and De Long \(1993\)](#). That paper shows how investors learning about the time-varying mean of a dividend process can lead to excess stock market movements; thus investors learning and updating about the time-varying mean of a dividend process can appear to "extrapolate" recent innovations in the dividend process. In contrast to the Barsky-DeLong framework, however, the mean of the demand factor in our framework is *not* time-varying (Equation (2)), but firms mistakenly *perceive* it to be (Equation (7)). In comparison to [Barsky and De Long \(1993\)](#), then, agents in our model end up "learning too much" from past demand shocks, leading to an excessive degree of extrapolation, and insufficient mean reversion, relative to the full-information benchmark, as pointed out above.

Dynamic investment decision. The model has three state variables at each point in time: A_t , \bar{A}_t , and Q_t . For an individual firm, its Bellman equation is

$$J(q_t; A_t, \bar{A}_t, Q_t) = \max_{i_t} \left\{ V(q_t, i_t; A_t, \bar{A}_t, Q_t) + \frac{\mathbb{E}_f[J(q_t + p \cdot i_t; A_{t+1}, \bar{A}_{t+1}, Q_{t+1} | A_t, \bar{A}_t, Q_t)]}{1 + r} \right\}. \quad (10)$$

Here " \mathbb{E}_f " means that the expectation is taken under firms' subjective beliefs. The first-order condition gives

$$P_r + k \cdot i_t^* = p \cdot \sum_j \frac{\mathbb{E}_f[\Pi_{t+j} | A_t, \bar{A}_t, Q_t]}{(1 + r)^j} \equiv p \cdot P(A_t, \bar{A}_t, Q). \quad (11)$$

Here $P(A_t, \bar{A}_t, Q_t)$ is a hypothetical price of discounting future expected per-unit net profit at the required rate of return r under firms' subjective expectations.

We now characterize the optimal level of investment in the proposition below.

Proposition 1. *In the investment model described above, firms' optimal level of investment is*

$$i_t^* = x + y_1 \cdot A_t + y_2 \cdot \bar{A}_t + z \cdot Q_t, \quad (12)$$

where

$$z = \frac{BMp^2 + kr}{2kp} - \sqrt{\left(\frac{BMp^2 + kr}{2kp}\right)^2 + \frac{BM}{k}}, \quad (13)$$

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} (1+r)k + BMp^2 - k\rho_f - kzp & -\rho_A\rho_fk \\ -(1-\rho_f)\alpha k & (1+r)k + BMp^2 - (1-\rho_f)\alpha\rho_Ak - k(1-\rho_A) - kzp \end{pmatrix}^{-1} \\ \times \begin{pmatrix} pM\rho_f \\ pM(1-\rho_f)\alpha \end{pmatrix}, \quad (14)$$

$$x = \frac{(ky_1 + \rho_Aky_2 + pM) \cdot (1-\rho_f)(1-\alpha)\bar{A} - rP_r - p(C + \delta P_r)}{k(r - zp) + BMp^2}. \quad (15)$$

Proof. See Appendix. ■

A. Impulse Response Analysis: the Cushioning Benefits

We now examine the model implications through some impulse response analyses. Figure 2 plots the net investment (new drilling) I_t , total wells Q_t , and the oil price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the model with biased beliefs ($\alpha = 0.9$ and $\rho_A = 0.25$); from the steady-state, a sequence of half standard deviation negative shocks on A_t are imposed at $t = 2, 3, 4, 5,$ and 6 .

The impulse responses presented in Figure 2 are computed using model parameters calibrated for the Alaskan oil exploration sector. (These parameter values are presented and discussed in Section III below.) These impulse responses highlight the cushioning benefits of biased beliefs. Compared to the benchmark case, biased beliefs lead to lower investments (new drilling) in the face of negative demand shocks, which lowers the number of active wells, and persists over many periods after the negative demand shocks are realized. This lower wellcount ‘‘cushions’’ the negative impact of the adverse demand shocks on the output price, resulting in a smaller price decline and a faster price recovery. To see the mechanism in more details, notice that after a sequence of negative shocks on A_t , overextrapolation leads firms to lower their estimation of the long-run mean of the demand factor \bar{A}_t^α , hence becoming pessimistic about future prices and therefore reducing their investment. Relative to the benchmark case, the number of wells and hence total oil production in subsequent periods drops to a larger degree and stays persistently low in the behavioral model. This comes from two reasons: first, firms' pessimistic beliefs about the future output price are persistent; second, lower past investments cumulatively result in lower total drilling in subsequent

periods. This supply effect of biased beliefs partially offsets the negative effect of adverse demand shocks on the output price. In this example, extrapolation reduces the decline of the output price by 8.2% and results in a faster price recovery at month 11 compared to month 15 for the benchmark case, a difference of 27%.

While this example has considered a sequence of negative demand shocks leading to a downturn, the cushioning effects also arise when the market is in an upturn. Indeed, an example with a sequence of positive demand shocks would produce exactly symmetric results: in the upturn, extrapolative firms would overinvest in new drilling projects, leading to excessive production accompanied by downward pressure prices. The cushioning effects here would be of the same magnitude, albeit of the opposite sign: extrapolation would reduce the magnitude of the price by 8.2% and prices would fall back to their pre-shock levels by month 11, four months earlier compared to the non-extrapolative benchmark case.

[Place Figure 2 about here]

Some additional observations are worth making. First, the cushioning benefits come at some cost, as large cutbacks in investment (a 27% reduction from 194 to 142 by period 8 in the middle panel of Figure 2) can imply a high level of industrial turmoil; this “shakeout” is more sizeable under extrapolative beliefs than in the non-extrapolative benchmark. As such, the overall welfare calculus of the cushioning effects is ambiguous, involving an intertemporal tradeoff between investment and output in the short run vs. faster recovery and higher output in the long run.

Second, the smaller decline of the output price in the behavioral model comes together with a faster recovery. As negative demand shocks continue to arrive at $t = 4$ and 5, they are basically offset by the lower drilling activity in the behavioral model. As a result, the output price stays relatively flat during these periods. Finally, combining a sequence of small negative demand shocks into a big shock tends to limit the cushioning effect. Figure 3 shows that if we clump all the half standard deviation negative demand shocks from Figure 2 into a large negative shock at $t = 2$, the decline of the output price is of the same magnitude in both the benchmark model and the behavioral model, although overextrapolation still leads to a faster recovery.

[Place Figure 3 about here]

B. Empirical Implications of Extrapolative Producers

In what follows, we study oil exploration in Alaska as an empirical application of the model. Since oil exploration is not a liquid asset market with ample trading and resale opportunities, the return regressions or survey evidence used in the existing literature to detect extrapolation are not available in this context. For that reason, before moving on to the empirical section of the paper, we derive some theoretical results to guide our empirical strategy of detecting extrapolative beliefs.

The following corollary shows how Proposition 1 pins down the coefficients of regressing current and future investments on the current output price.

Corollary 1. *The regression coefficient for regressing I_t on H_t , both conditional and unconditional on A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , is $\beta_0 = y_1 + \rho_A y_2$. The regression coefficient for regressing I_{t+1} on H_t , both conditional and unconditional on A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , is $\beta_1 = (y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A + zp(y_1 + \rho_A y_2)$.*

Proof. See Appendix. ■

We plot in Figure 4 the coefficients β_0 and β_1 as functions of α and ρ_A . Compared to the benchmark case of $\alpha = 0$ or $\rho_A = 0$, higher values of α and ρ_A make firms tend to overestimate the long-run mean of the demand factor after a sequence of positive demand shocks, hence reducing firms' perceived degree of mean reversion about future prices and causing them to overinvest; overinvestment after a sequence of high prices — high prices are caused by positive demand shocks — therefore leads to higher values of β_0 and β_1 . Note that β_1 tends to be lower than β_0 as firms still anticipate some degree of mean reversion about future prices so they tend to scale back future investment. Also note that as ρ_A further increases, β_1 increases at a lower rate or decreases in some other cases: a higher ρ_A makes the estimated long-run mean of the demand factor \bar{A}_t^α less persistent, so adverse demand shocks in the future tend to affect investment decisions to a larger extent.

[Place Figures 4 and 5 about here]

Finally, we consider some implications of our model for firms' investment after different price pattern scenarios. In Figure 5 we plot the impulse responses for our model under two different price pattern scenarios. From the steady-state, two different sequences of shocks on A_t are imposed from $t = 2$ to $t = 6$, resulting in one case a steady rise in oil price and in another case an increasing (or accelerating) rise in price. The top panel of the figure then suggests that firms overinvest less (by 8.9%) in the accelerating price scenario compared to the steady price increase scenario. Indeed, our model suggests that, in comparison with a steady increase in price, an accelerating price increase leads to smaller revisions in the firms' estimated long-run mean \bar{A}_t^α , giving rise to stronger perceived mean reversion in future prices and less overinvestment.

It is worth noting that leading asset pricing models of extrapolation such as Barberis, Shleifer, and Vishny (1998) and Barberis et al. (2015) cannot explain this finding: in these models, an increasing rise will result in stronger extrapolative beliefs and therefore more investment. On the other hand, our finding is consistent with the empirical results of Barber, Odean, and Zhu (2009) and Greenwood, Shleifer, and You (2017) in the context of the stock market. These findings suggest some reasonable caution on the part of extrapolators in the face of accelerating price increase, and hence a belief formation process that is more sophisticated than simple extrapolation studied in the

literature. To the best of our knowledge, the belief dynamic we proposed in equations (7) and (8) is among the first that can, at least in part, capture this path-dependent feature of belief formation.⁷

II. Oil Exploration: An Empirical Application

Having established the existence of cushioning benefits from extrapolative beliefs in a theoretical model, we now proceed to quantify them in a real world setting. In choosing a suitable industry for a case study, we seek an industry with pronounced boom and bust periods, where cushioning can play a more important role. For that reason, we focus on oil exploration and production in Alaska. There are several reasons which make Alaska a suitable stage for our analysis: the size of its market, the purity of the North Slope oil price index, and the ability for Alaskan operators to respond to price changes. Alaska has long been one of the biggest oil producing states in the U.S, consistently ranked amongst the top five across oil producing regions in the country since the 1970s. Its production peaked at around two million barrels a day in 1988, behind only Texas. In addition to the sheer size of the production in Alaska, petroleum activities in Alaska are centered in the remote North Slope region (along the northern Arctic coast of Alaska). North Slope field production has accounted for over 95% of all field production in Alaska each year since the latter half of 1970s (EIA (2017)). As a result, the Alaska North Slope crude oil price is a pure indicator of the supply and demand for Alaska-produced oil, as almost no supply from regions outside of Alaska is in the mix. Furthermore, due to its Arctic location, the Alaska North Slope region has a topography and geology which differ from other main oil-producing regions in the continental United States, such as California and Texas, and which pose unique challenges for oil drilling and production. Successful drilling in this region requires a fair amount of local know-how and specialization. Due to such reasons and also different state regulations, international large operators, such as British Petroleum and ConocoPhillips, all have independent operations in Alaska that are separate from their other U.S. operations. BP, for instance, has an independently incorporated subsidiary in Alaska, BP Exploration (Alaska) Inc (BP (2016)). This independence allows these subsidiaries to make drilling decisions themselves, and it is thus reasonable to model their decisions in Alaska as independent from the decisions in other regions. These independent operations are also arguably more responsive to the boom and bust cycle in the Alaska oil industry. For that reason, the cushioning effects that we pointed out in the previous section, which soften out the boom and the busts, may be especially important in Alaska compared to other oil-producing regions.

Before proceeding to the calibration exercises which quantify the magnitude of the cushioning benefits for several historical episodes in Alaskan oil exploration, we present some suggestive evidence from regressions which indicate that oil producers are affected by extrapolative beliefs in their oil production decisions. In these regressions, we are guided by the empirical implications from Corollary 1 above.

⁷Additional computations also showed that the biased belief specification for ship producers in Greenwood and Hanson (2015) cannot generate this pattern of decreasing overinvestment with an accelerating price pattern.

A. Data Description

The data used in our case study is drawn from multiple sources. For oil prices, we use the monthly Alaska North Slope (ANS) first purchase price per barrel from the U.S. Energy Information Administration (EIA) from 1977 to 2016. As depicted by the first panel of Figure 6, the ANS price remained relatively low prior to 2000, but increased drastically after 2000. In the meantime, cost to drill an oil well, using the nominal average cost per crude oil well drilled each year in the US provided by the EIA, followed a similar trend. The drilling cost remained relatively stable but skyrocketed after 2000, as shown in the middle panel of Figure 6.⁸ To more accurately capture firms' revenue and cost considerations, we introduce a *normalized oil price*, defined by ANS oil price divided by the oil well cost for the corresponding month, and then multiplied by the drilling cost of June, 2000. Hence the normalized price in June 2000 is the same as the nominal ANS price, but the normalized oil price is higher in the 1980s and much lower in the 2000s, as shown in the right panel in Figure 6.

[Place Figure 6 about here]

To examine how firm investments respond to past levels of oil prices, we focus on the number of wells drilled for oil producing or servicing purposes, using historical well activity data from the Alaska Oil and Gas Conservation Commission (AOGCC). The drilled wells are those with records showing the dates of actual construction activities, such as well spudding or well completion, or those with positive well depth. A firm with extrapolative beliefs is more likely to drill a well when it observes high oil prices in the recent past because the firm perceives high prices and hence high revenue moving forward. In Figure 7, we plot the time series of the number of wells drilled in each month, together with the normalized oil price. It suggests that indeed a larger number of wells are drilled following high price periods.

[Place Figure 7 about here]

While Figure 7 suggests that recent price levels seem to be associated with firms' well drilling decisions, these decisions may be fully rational. Rather than over-extrapolating, the firm may correctly foresee high oil prices in the future after observing high oil prices in the recent past, and therefore wells drilled at these high price level times create more profit due to the rising price. Additional evidence is needed before concluding that firms over-responded to the recent high prices, in line with the over-extrapolation hypothesis. In Figure 8, we plot, on a monthly basis, the time series of the average per-well profit and production in the first 60 months of production for wells that are drilled, along with the normalized oil prices.⁹ Both plots do *not* seem to support that the

⁸The cost data is only available on an annual basis from 1960 to 2007. We use linear interpolation to obtain monthly drilling cost data.

⁹Profit is calculated using normalized price multiplied by the production. As the normalized price is a measure of how profitable oil production is, we multiply it by production volume for total profit

firms are rationally foreseeing future oil prices. On the contrary, both plots in Figure 8 suggest that *lower-revenue and less productive* wells are drilled following periods of high oil prices, more consistent with the interpretation of these wells as mistakes arising from over-extrapolative beliefs.

[Place Figure 8 about here]

To further examine firms' extrapolation bias, we utilize the availability of well approval data to study the percentage of scrapped wells in each month. Within the well history data, we observe all wells that received permits to drill. However, not all permitted wells ended up being drilled. Scrapped wells are those that received permits but the constructions of which never took place within 24 months of the permit approval.¹⁰ As stipulated by the Alaska State Legislature (Title 20 Chapter 25), if a well is not drilled within this period, a new permit needs to be applied for. If a firm has extrapolative beliefs, then it might become overly exuberant and apply for more well permits after observing high oil prices, but subsequently reverse its investment decision when price drops after the permits are issued. To capture this reversal of investment decisions, we look at the proportion of wells scrapped each month, calculated as the number of wells ending up scrapped over the total number of wells approved in each month. Figure 9 plots the time series of wells scrapped percentage, as well as the normalized oil price. It suggests that the percentage of wells scrapped increases following periods of recent low prices.

[Place Figure 9 about here]

B. Regression Results

Based on what we observe in Figures 7, 8, and 9, we now run a series of regressions to formally test whether firms extrapolate recent oil prices when making investment decisions, whether this extrapolative belief is unbiased in terms of well profit and production, and whether firms take actions to correct their investment decisions after receiving drilling permits.

First we look at the relationship between the number of wells drilled each month and normalized oil prices lagged by various time periods: one, three, six, nine and twelve months,

$$Y_t = \alpha + \beta P_{t-\tau} + \gamma S_t + \varepsilon_t, \quad \text{where } \tau = 1, 3, 6, 9 \text{ or } 12 \text{ months.} \quad (16)$$

Here, Y_t is the number of wells drilled in month t ; $P_{t-\tau}$ is the normalized oil price from τ periods before; and ε_t is the error term for each period. In this equation, we also include price pattern variable S_t , an indicator variable for whether prices have been changing at an increasing or decreasing rate in the previous months leading up to month t . We categorize price paths into five

¹⁰However, if a different well was drilled at the exact same longitude and latitude within two years of the initial approval, the initial well does not count as a scrapped well, since a better well may be planned to replace the current one.

patterns: no clear pattern in the rate of change as the reference level; decreasing rate of decline; increasing rate of decline; decreasing rate of rise; and increasing rate of rise. These categories are calculated as follows. Let n be the maximum number of consecutive months prior to t that prices have gone up or down. For n that is greater than or equal to three, we calculate the rate of change as the difference of price changes between the second and first half of these periods.¹¹ If the rate of change has a different sign from the price trend, then this period experiences decreasing rate of change; otherwise it experiences increasing rate of change. For instance, if the rate of change is less than zero but the price has been going up, then this period is going through a decreasing rise in price.

With extrapolative beliefs, firms perceive high oil prices moving forward after observing high recent prices, resulting in high investments in the current period. In other words, extrapolation implies a positive relation between the number of wells drilled and past oil prices; coefficient β in (16) is expected to be positive. As for the price patterns, some empirical results from previous papers for financial markets suggest that γ should be negative when S_t corresponds to an increasing rate of rise; that is, accelerations in the rate of price increase tend to reduce extrapolative effect of past prices on current behavior (see Barber et al. (2009) and Greenwood et al. (2017)). It is an empirical question, however, whether such a result will obtain outside financial markets. Also note that, as we described earlier, traditional extrapolation models cannot capture this fact, whereas the belief dynamic proposed in equations (7) and (8) tends to generate it.

Table I summarizes the regression results based on equation (16). In this table, we see a significant positive relation between the monthly number of wells drilled and the lagged normalized oil prices from six months or more prior to drilling; the relation becomes insignificant when the lag of past oil price is less than six months. This reduced significance for the more recent months' coefficients is perhaps due to two reasons. First, it takes time to move equipment and set up rigs before drilling can actually start. Second, it takes time for firms to attend to past oil prices; this interpretation with limited attention is consistent with models of extrapolation in the behavioral finance literature (see, for instance, Barberis et al. (2017), Barberis and Shleifer (2003), and Hong and Stein (1999)). Overall, Table I suggests that oil companies in Alaska increase their investment in well drilling when the oil price levels are high in recent months.

[Place Table I about here]

In addition, we also find that the extrapolation is tempered by the pattern of increasing rise in price, as the coefficient on the “increasing rise” price patten is negative and statistically significant. That is, when price has been going up at an increasing rate, all else equal firms become less extrapolative and therefore reduce investments. This confirms the predictions of the extrapolation models illustrated in Figure 5. For instance, if the normalized oil price is \$50 six months ago with

¹¹If n is even, the difference is $(P_t - P_{t-n/2}) - (P_{t-n/2} - P_{t-n})$. If n is odd, the difference is $[P_t - (P_{t-(n-1)/2} + P_{t-(n+1)/2})/2] - [(P_{t-(n-1)/2} + P_{t-(n+1)/2})/2 - P_{t-n}]$, which can be simplified to $(P_t - P_{t-(n-1)/2}) - (P_{t-(n+1)/2} - P_{t-n})$.

no observable price pattern, then a company with extrapolative beliefs would drill 18.1 wells in the current month, 1.1 more than if the price was \$40 six months ago. Furthermore, if the price has been rising at an increasing rate in the months leading up to the current month, then the number of wells drilled would become 16, which is 12% lower than the original level of 18.1. In other words, the impact of observing increasingly rising oil prices leading up to the month of drilling is similar to a decrease of \$20 in oil prices from six months prior.

[Place Tables II and III about here]

The previous regression provides some evidence on firms extrapolating high levels of past prices that they observe. However, firms' extrapolative beliefs can be fully rational; empirically, this regression is not able to disentangle over-extrapolative drilling behavior from "rational exuberance" when recent oil prices have been high. To address this issue, we run regressions of initial five-year profit and production of wells on lagged normalized prices one, three, six, nine, and twelve months prior to the drilling date, controlling for price trend patterns. The initial five-year production is calculated as the amount of oil produced in the first 60 months of the well production. The initial five-year profit is the monthly production multiplied by monthly normalized oil price in the first 60 months of well production. Tables II and III summarize the results for these two regressions, which confirms the graphical evidence presented in Figure 8 which we discussed earlier. Table II shows that well profits are actually *lower* for wells drilled following high price levels, suggesting that firms over-extrapolate and fail to foresee the price reversal when they observe periods of high prices. Similarly, Table III shows that wells drilled following high price levels are not more productive, hence ruling out the possibility that firms save the most productive wells for high price periods. Taken together, these results support the interpretation that firms drill *excessively* after a period of high prices, leading to lower profit and production, and symptomatic of biased beliefs.

[Place Table IV about here]

Just as how firms can over-extrapolate when they drill a well, firms can also over-extrapolate when they apply for the permit to drill. When they observe oil price drops after the approval, they can retract their initial plan and decide not to drill. Hence, observing an increased likelihood of scrapped wells when prices drop after initial approval can be evidence of firms being overly exuberant when they make drilling plans in the first place. To test this hypothesis, we regress the indicator of whether a well is scrapped on the change in oil prices 24 months after the permit issuance, controlling for the current normalized price in the approval month. Table IV summarizes the results. The changes in oil prices in both columns are calculated as the percentage difference between average oil prices in the next 24 months and the price in the approval month. In Column 2, we also control for additional well heterogeneity, such as the region of the wells, the operator and the unit that the wells belong to.

Qualitatively, the results in both columns agree, and we will focus on the results in Column 2.

Starting from the top, the negative and significant coefficient on the normalized oil price (-0.001) shows that when the oil price levels are high at the time of approval, firms are less likely to scrap these wells. This observation confirms the earlier results from Table I that when oil prices are high, firms are more likely to actually drill the wells, and are less likely to scrap these wells.

The negative and significant coefficient (-0.009) on the 2-year post-approval price change shows that an increase in the price after a well has been approved lowers the probability of scrapping a well. This result is consistent with extrapolative producers, as extrapolative firms, which may be eager to initiate new drilling projects when prices have been high (as demonstrated in earlier results) end up scrapping these projects if oil prices decline after obtaining drilling permits (but before drilling has begun).

III. Model Calibration

Having presented empirical evidence consistent with the hypothesis of extrapolation in the oil exploration sector, we now turn to the calibration of the theoretical model presented in Section I using data from Alaska to set model parameters.

To more realistically measure the effect of firms' extrapolation bias, we calibrate the key parameters in the model using the available Alaska data. Table V summarizes the list of parameters, their values and our justification for choosing each value.

[Place Table V about here]

Demand-related parameters are associated with equations (1) and (2). A_t in equation (1) can be viewed as the long-term mean \bar{A} plus a stochastic term specific to period t . Hence, to obtain \bar{A} and B , we regress the monthly per-barrel North Slope price, H_t , on the number of productive wells in Alaska in each month, Q_t .¹² We will refer to this regression as the demand regression. The resulting intercept from the demand regression is then \bar{A} and the coefficient associated with well count is then $-B$. The result of the demand regression can be found in Table VI.¹³ Next, in equation (2), notice that $A_{t+1} - \bar{A}$ is simply the residual from the demand regression for period $t+1$. Thus, to obtain ρ_0 , we regress the residual on its one-month lag, and to obtain σ_ε , we calculate the standard deviation of the residuals from this lag regression. The results of these regressions can be found in Table VII.¹⁴

[Place Tables VI and VII about here]

¹²Being "productive" here is defined as a monthly production of at least 2000 barrels, around the median of what the bottom 10% of wells produce monthly in our sample.

¹³For robustness, we ran simple linear regression as well as IV regressions using lagged well counts as the instruments. All regressions give us comparable range of parameter values.

¹⁴The estimate of $\rho_0 = 0.68$ indicate a substantial degree of mean reversion in the process of A_t , the demand factor. Note that this does not rule out the possibility that oil *prices* may not be mean reverting.

To determine the representative firm’s investment level, as shown in equation (3), we need to estimate the depreciation rate δ and the probability of successful drilling p . For the depreciation rate, we look at all the development wells that were drilled within our time frame, and see for how long they produce. The median length of production life for our wells is around 180 months, and hence the depreciation rate is on average $1/180$. The success rate amongst wells for production purposes, exploratory and development wells, varies by the definition of success. If the definition of success is producing over 1000 barrels, then the success rate is around 80%.¹⁵

To calculate a representative firm’s profit, we need well production in a month for an representative well, M , and operating cost C , as well as drilling-related cost parameters P_r and k . We look at the median well for each month since July, 1978, and find that the average of all the months come to around 13,000 barrels. The operating costs include lease operating expenses, gathering, processing and transporting costs, as well as water disposal and General and Administrative costs. These costs vary widely across well locations, performances or the amount of production (EIA (2016)). We use C as an umbrella term for all of these costs, and set C to be around 100,000 a month as an estimate for the monthly level of all costs mentioned above. Finally, drilling cost-related parameters can be extracted from the well cost data. In equation (6), the cost per well is modeled as $P_r + k \cdot i_t/2$, which increases as the number of wells drilled increases. Knowing the total cost of drilling per well and the number of newly-drilled well in each month, i_t , we regress the monthly well drilling cost on the number of wells drilled each month, and the intercept from the regression is P_r and the coefficient for the number of wells is $k/2$. The results of the cost regression can be found in Table VIII.

[Place Table VIII about here]

The impulse responses emerging from the model simulations utilizing these parameters are in Figures 2 and 3, and we already discussed them earlier.

A. *Two Historical Episodes: 2008-2011 and 1986-1987*

Taking the calibration exercise one step further, we next present two additional model-fitting exercise to quantify the cushioning benefits for specific historical episodes of price downturns for Alaskan crude oil. Since oil produced in Alaska is mostly supplied to other parts of the U.S. and can be easily substituted with oil produced from other regions in the U.S. and areas outside of the U.S., Alaska oil prices are sensitive to domestic and international events that affect oil demand and supply. We focus on episodes following two such events, the U.S. financial crisis in 2008 (the Great Recession) and Saudi Arabia’s dramatic increase in production in 1986. These are illustrated in Figures 10 and 11.

¹⁵If we make a more stringent requirement that the development well production needs to be at least 1 million barrels, then $p = 0.55$. For a threshold of 0.5 million, $p = 0.65$; for a threshold of 100,000 barrels, $p = 0.75$.

[Place Figure 10 about here]

In Figure 10, we consider a forty-month period from August 2008 to November 2011, coinciding with the most recent financial crisis in the US. As illustrated in the bottom panel of this figure (in the bubble-dashed line), oil prices fell sharply during the first six months (08/2008-01/2009) by over 60% from peak to trough, and then recovered very slowly and gradually, regaining the initial price level at month 33 (04/2011). Based on this, we calibrated a shock process for the extrapolative model (dashed green line) to match the relative magnitudes and shape of the actual price process. The recovery for this calibrated extrapolative price process takes 33 months, as in the actual price process. Using this calibrated shock process, we also simulated prices in the benchmark non-extrapolative industry (graphed in solid blue). For this benchmark economy, we see that the recovery takes longer; only at month 37 (08/2011) do prices reach the initial level. Thus, in this way, we find that the cushioning benefits shortened the recovery process by roughly 4 months, or 11%.

The two top panels in the figure show striking differences in the extrapolative and benchmark industry during the price drop and recovery process, which illustrates the complex welfare effects of biased beliefs. With extrapolation, investment falls sharply as prices drop, leading to large differences in the well-count. Since the oil exploration sector is composed of many small firms, such a large drop in drilling activity will entail a sizeable “shakeout” as firms become inactive and are forced to leave the market. We can get a sense of the size of this “shakeout” by looking at a period with similar drastic decline of oil prices, though to a smaller extent and without the woes of the global financial crisis. Following the oil price collapsing by 40% in the second half of 2014, 128 oil and gas companies filed for bankruptcy between 2015 and 2016, up from around 30 between 2013 and 2014 (Haynes and Boone (2016), Egan (2016)). At 30 months, indeed, the number of active wells in the extrapolative industry is only a fraction of the well-count in the benchmark industry. Clearly, the accompanying decrease in output allows prices to recover more quickly. Extrapolative beliefs lead to a direr but shorter duration, and results in ambiguous welfare effects.

[Place Figure 11 about here]

Figure 11 illustrates a similar exercise for an earlier and less pronounced price drop episode, in the mid-1980's. Between 1981 and 1985, Saudi Arabia reduced its oil production by three quarters in order to combat the price collapse caused by the world consumption decline. However, beginning in 1986, Saudi Arabia decided to abandon its effort and ramp up its production, causing oil prices to fall further in 1986 (Hamilton (2011)). We again compare the lengths of recovery period following this event under the behavioral and benchmark models. By design, the recovery occurred in the fifteenth month (April 1987) in the actual data, and also for the behavioral model. For the benchmark non-extrapolative model, however, the recovery did not occur until September 1987; in this example, then, the cushioning effects shorten the recovery process by five months.

Naturally, these are stylized examples, but they illustrate the real benefits that “cushioning” can have in a real-world setting, in an industry notorious for its boom and busts sequences. Indeed, a lesson from these examples is that the cushioning benefits of extrapolation can soften the extremes of the cycles.

IV. Conclusion

Much of the existing literature in behavioral economics and finance has focused on the negative and undesirable effects of behavioral biases and biased beliefs. In contrast, we point out in this paper that in certain settings, such as industries prone to periodic boom and bust cycles, biased beliefs can have benefits in terms of softening the up and downs of the economic cycle. In these industries, biased beliefs cause firms making investment decisions to respond more quickly to recent information in market prices. Thus, for instance, a price downturn will trigger a more immediate decrease in investment; in turn, this leads to lower supply which “cushions” and prevents prices from falling too quickly and leads to a quicker recovery. Modelling and quantifying these positive implications of biased beliefs on industry dynamics are important contributions of this paper.

We develop a theoretical framework, based on a standard aggregate investment model to illustrate these cushioning benefits. We then apply this model to the oil exploration industry in Alaska, a highly volatile industry characterized by sharp price fluctuations. One striking calibration example shows that the industry downturn during the 2008 financial crisis would have been lengthened by four months if oil industry firms did *not* have biased extrapolative beliefs. This suggests that the cushioning benefits can be sizable in a real-world setting. In ongoing work, we are exploring other sectors in which extrapolative beliefs may be important and hence in which the cushioning benefits can play an important role.

Appendices

A. Analytical Results for the Model

Proof of Proposition 1. We conjecture and verify later that the optimal investment is linear in state variables A_t , \bar{A}_t , and Q_t

$$i_t^* = x + y_1 \cdot A_t + y_2 \cdot \bar{A}_t + z \cdot Q_t. \quad (\text{A.1})$$

Equation (11) then implies

$$P(A_t, \bar{A}_t, Q_t) = (kx + ky_1 \cdot A_t + ky_2 \cdot \bar{A}_t + kz \cdot Q_t + P_r)/p. \quad (\text{A.2})$$

By applying the law of iterated expectations on (11), firms derive

$$P_r + k \cdot i_t^* = p \cdot \frac{\mathbb{E}_f[\Pi_{t+1} + P(A_{t+1}, \bar{A}_{t+1}, Q_{t+1}) | A_t, \bar{A}_t, Q_t]}{1+r}. \quad (\text{A.3})$$

Equations (4), (5), (7), and (8) allow us to write (A.3) out as

$$\begin{aligned} & P_r + kx + ky_1 \cdot A_t + ky_2 \cdot \bar{A}_t + kz \cdot Q_t \\ = & p \cdot \frac{M\{\alpha\bar{A}_t + (1-\alpha)\bar{A} + \rho_f[A_t - \alpha\bar{A}_t - (1-\alpha)\bar{A}] - B(Q_t + px + py_1 \cdot A_t + py_2 \cdot \bar{A}_t + pz \cdot Q_t)\} - C - \delta P_r}{1+r} \\ & + \frac{kx + (ky_1 + \rho_A ky_2) \cdot \{\alpha\bar{A}_t + (1-\alpha)\bar{A} + \rho_f[A_t - \alpha\bar{A}_t - (1-\alpha)\bar{A}]\} + ky_2 \cdot (1-\rho_A)\bar{A}_t}{1+r} \\ & + \frac{kz \cdot (Q_t + px + py_1 \cdot A_t + py_2 \cdot \bar{A}_t + pz \cdot Q_t) + P_r}{1+r}. \end{aligned} \quad (\text{A.4})$$

The fact that both sides of (A.4) are linear functions of A_t , \bar{A}_t , and Q_t verifies the conjecture in (A.1). Matching terms in a sequential order then solves for x , y_1 , y_2 , and z . First, matching terms for Q_t gives the solution of z in (13). Then matching terms for A_t and \bar{A}_t , we obtain

$$\begin{aligned} ky_1 &= pM \frac{\rho_f - Bpy_1}{1+r} + \frac{(ky_1 + \rho_A ky_2)\rho_f + kzpy_1}{1+r}, \\ ky_2 &= pM \frac{(1-\rho_f)\alpha - Bpy_2}{1+r} + \frac{(ky_1 + \rho_A ky_2)(1-\rho_f)\alpha + ky_2 \cdot (1-\rho_A) + kzpy_2}{1+r}. \end{aligned} \quad (\text{A.5})$$

Notice that y_1 and y_2 are interrelated because the evolution of \bar{A}_t is driven by past realizations of A_t . Solving these two simultaneous equations then leads to (14). Finally, matching the constant term gives (15). ■

Proof of Corollary 1. Conditional on knowing A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , I_{t-1} and therefore Q_t are

both determined. In this case, the movements of H_t and I_t are only caused by the realization of the random shock ε_t . That is

$$\begin{aligned}
I_t &= x + (y_1 + \rho_A y_2) \cdot H_t + y_2(1 - \rho_A) \cdot \bar{A}_{t-1} + (z + y_1 B + \rho_A y_2 B) \cdot Q_t \\
&= x + (y_1 + \rho_A y_2) \cdot H_t + y_2(1 - \rho_A) \cdot \bar{A}_{t-1} + (z + y_1 B + \rho_A y_2 B) \cdot f(A_{t-1}, \bar{A}_{t-1}, Q_{t-1}).
\end{aligned} \tag{A.6}$$

So the coefficient for regressing I_t on H_t , both conditional and unconditional on A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , is $\beta_0 = y_1 + \rho_A y_2$.

We now consider the coefficient of regressing I_{t+1} on H_t . Conditional on A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , the realization of ε_t determines H_t and I_t , which further determine Q_{t+1} . Then the realization of ε_{t+1} determines A_{t+1} , \bar{A}_{t+1} , and I_{t+1}

$$\begin{aligned}
I_{t+1} &= x + y_1 \cdot A_{t+1} + y_2 \cdot \bar{A}_{t+1} + z \cdot Q_{t+1} \\
&= x + (y_1 + \rho_A y_2) \cdot [\bar{A} + \rho_0(A_t - \bar{A}) + \varepsilon_{t+1}] + y_2(1 - \rho_A)\bar{A}_t + z \cdot [Q_t + pI_t] \\
&= x + (y_1 + \rho_A y_2)\varepsilon_{t+1} + (y_1 + \rho_A y_2)(1 - \rho_0)\bar{A} + [(y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A]H_t \\
&\quad + y_2(1 - \rho_A)^2\bar{A}_{t-1} + \{z + B[(y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A]\}f(A_{t-1}, \bar{A}_{t-1}, Q_{t-1}) \\
&\quad + zp[x + (y_1 + \rho_A y_2)H_t + y_2(1 - \rho_A)\bar{A}_{t-1} + (z + y_1 B + \rho_A y_2 B)f(A_{t-1}, \bar{A}_{t-1}, Q_{t-1})] \\
&= x + (y_1 + \rho_A y_2)\varepsilon_{t+1} + (y_1 + \rho_A y_2)(1 - \rho_0)\bar{A} \\
&\quad + [(y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A + zp(y_1 + \rho_A y_2)]H_t \\
&\quad + y_2(1 - \rho_A)^2\bar{A}_{t-1} + \{z + B[(y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A]\}f(A_{t-1}, \bar{A}_{t-1}, Q_{t-1}) \\
&\quad + zp[x + y_2(1 - \rho_A)\bar{A}_{t-1} + (z + y_1 B + \rho_A y_2 B)f(A_{t-1}, \bar{A}_{t-1}, Q_{t-1})].
\end{aligned} \tag{A.7}$$

So the coefficient for regressing I_{t+1} on H_t , both conditional and unconditional on A_{t-1} , \bar{A}_{t-1} , and Q_{t-1} , is $\beta_1 = (y_1 + \rho_A y_2)\rho_0 + y_2(1 - \rho_A)\rho_A + zp(y_1 + \rho_A y_2)$. ■

REFERENCES

- Abel, Andrew B., 1981, A dynamic model of investment and capacity utilization, *Quarterly Journal of Economics* 96, 379–403.
- Abel, Andrew B., and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 84, 1369–1384.
- Alti, Aydogan, and Paul Tetlock, 2014, Biased beliefs, asset prices, and investment: A structural approach, *Journal of Finance* 69, 325–361.
- Amromin, Gene, and Steven A. Sharpe, 2013, Expectation of risk and return among household investors: Are their sharpe ratios countercyclical?, Working paper, Federal Reserve Board, Washington DC.
- Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop, 2009, Predictability in financial markets: What does survey expectations tell us?, *Journal of International Money and Finance* 28, 406–426.
- Bagehot, Walter, 1873, *Lombard Street: A Description of the Money Market* (Wiley, New York).
- Barber, Brad M., Terrance Odean, and Ning Zhu, 2009, Systematic noise, *Journal of Financial Markets* 12, 547–569.
- 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, 2017, 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.

- Barsky, Robert B, and J Bradford De Long, 1993, Why does the stock market fluctuate?, *Quarterly Journal of Economics* 108, 291–311.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2017, Diagnostic expectations and credit cycles, *Journal of Finance* forthcoming.
- BP, 2016, BP Annual Report and Form 20-F 2016 <https://www.bp.com/content/dam/bp/en/corporate/pdf/investors/bp-annual-report-and-form-20f-2016.pdf>.
- Cassella, Stefano, and Huseyin Gulen, 2015, Extrapolation bias and the predictability of stock returns by price-scaled variables, Working paper, Purdue University.
- Choi, James, and Thomas Mertens, 2013, Extrapolative expectations and the equity premium, Working paper, New York University and Yale University.
- Conlin, Michael, Ted O’Donoghue, and Timothy J Vogelsang, 2007, Projection bias in catalog orders, *American Economic Review* 97, 1217–1249.
- Dong, Ming, David Hirshleifer, and Siew Hong Teoh, 2017, Stock market overvaluation, moon shots, and corporate innovation, Working paper, University of California, Irvine.
- Egan, Matt, 2016, U.S. oil bankruptcies spike 379% <http://money.cnn.com/2016/02/11/investing/oil-prices-bankruptcies-spike/index.html>.
- EIA, 2016, Trends in U.S. oil and natural gas upstream costs <https://www.eia.gov/analysis/studies/drilling/pdf/upstream.pdf>.
- EIA, 2017, Crude oil production, https://www.eia.gov/dnav/pet/pet_crd_crdpdn_adc_mbbldpd_a.htm.
- Fuster, Andreas, Benjamin Hebert, and David Laibson, 2011, Natural expectations, macroeconomic dynamics, and asset pricing, *NBER Macroeconomics Annual* 26, 1–48.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer, 2015, Expectations and investment, *NBER Macroeconomics Annual* 30, 379–442.
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2012, Neglected risks, financial innovation, and financial fragility, *Journal of Financial Economics* 104, 452–468.

- Glaeser, Edward L., Joseph Gyourko, and Albert Saiz, 2008, Housing supply and housing bubbles, *Journal of Urban Economics* 64, 198–217.
- Glaeser, Edward L., and Charles G. Nathanson, 2015, An extrapolative model of house price dynamics, Working paper, Harvard University and Northwestern University.
- Greenwood, Robin, Sam Hanson, and Lawrence Jin, 2016, A model of credit market sentiment, Working paper, California Institute of Technology and Harvard University.
- Greenwood, Robin, and Samuel Hanson, 2015, Waves in ship prices and investment, *Quarterly Journal of Economics* 130, 55–109.
- Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.
- Greenwood, Robin, Andrei Shleifer, and Yang You, 2017, Bubbles for fama, Working paper, Harvard University.
- Hamilton, James D., 2011, Historical oil shocks, *NBER Working Paper Series* 16790.
- Haynes, and LLP Boone, 2016, Oil patch bankruptcy monitor http://www.haynesboone.com/~media/files/attorney%20publications/2016/energy_bankruptcy_monitor/oil_patch_bankruptcy_20160106.ashx.
- 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 C. Stein, 1999, A unified theory of underreaction, momentum trading and overreaction in asset markets, *Journal of Finance* 55, 2143–2184.
- Jin, Lawrence J., 2015, A speculative asset pricing model of financial instability, Working paper, California Institute of Technology.
- Jin, Lawrence J., and Pengfei Sui, 2017, Asset pricing with return extrapolation, Working paper, California Institute of Technology.
- Kellogg, Ryan, 2014, The effect of uncertainty on investment: evidence from texas oil drilling, *American Economic Review* 104, 1698–1734.

Kindleberger, Charles P., 1978, *Manias, Panics, and Crashes: A History of Financial Crises* (Basic Book, New York).

Koijen, Ralph S. J., Maik Schmeling, and Evert B. Vrugt, 2015, Survey expectations of returns and asset pricing puzzles, Working paper, London Business School.

Kuchler, Theresa, and Basit Zafar, 2016, Personal experiences and expectations about aggregate outcomes, Working paper, New York University and Federal Reserve Bank of New York.

Vissing-Jorgensen, Annette, 2004, Perspectives on behavioral finance: Does “irrationality” disappear with wealth? Evidence from expectations and actions, *NBER Macroeconomics Annual 2003* 18, 139–194.




Figure 1. Illustration of the Biased Belief Process and Insufficient Mean Reversion. This figure contrasts the true process for the demand shocks (Equation (2)) from producers' beliefs (Equations (7) and (8)). The process of demand factor A_t is plotted in the solid line. The true long-run mean, \bar{A} , is invariant over time and equal to zero in this example, as plotted in the dotted line. However, producers' beliefs about the long-run mean, \bar{A}_t^α , change over time, and are plotted in the dash-dot line. The difference between A_t and \bar{A}_t^α measures the degree of mean reversion in A_t perceived by producers. The parameter values used in this example are: $\bar{A} = 0$, $\rho_0 = 0.68$, $\rho_f = 0.68$, $\sigma_\varepsilon = 4.25$, $\alpha = 0.9$, and $\rho_A = 0.25$.




Figure 2. Impulse Responses for the Benchmark Model and the Behavioral Model.

From the steady-state, a sequence of half standard deviation negative shocks on A_t are imposed at $t = 2, 3, 4, 5,$ and 6 . We plot the net investment I_t , total production Q_t , and the output price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the behavioral model ($\alpha = 0.9$ and $\rho_A = 0.25$). The other parameter values are: $B = 0.02, \bar{A} = 12, \delta = 0.6\%, r = 0.5\%, k = 22.8, C = 100, P_r = 463, \sigma_\varepsilon = 4.25, p = 0.8, \rho_0 = 0.68, \rho_f = 0.68,$ and $M = 13$.




Figure 3. Impulse Responses for the Benchmark Model and the Behavioral Model.

From the steady-state, a two and a half standard deviation negative shock on A_t is imposed at $t = 2$. We plot the net investment I_t , total production Q_t , and the output price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the behavioral model ($\alpha = 0.9$ and $\rho_A = 0.25$). The other parameter values are: $B = 0.02$, $\bar{A} = 12$, $\delta = 0.6\%$, $r = 0.5\%$, $k = 22.8$, $C = 100$, $P_r = 463$, $\sigma_\varepsilon = 4.25$, $p = 0.8$, $\rho_0 = 0.68$, $\rho_f = 0.68$, and $M = 13$.




Figure 4. Coefficients of Regressing Current and Future Investments on Current Price Level. The figure plots the coefficients of regressing current and future investments on the current price level, β_0 and β_1 , as functions of the belief-based parameters α and ρ_A . The default values for α and ρ_A are 0.9 and 0.25, respectively. The other parameter values are: $B = 0.02$, $\bar{A} = 12$, $\delta = 0.6\%$, $r = 0.5\%$, $k = 22.8$, $C = 100$, $P_r = 463$, $\sigma_\varepsilon = 4.25$, $p = 0.8$, $\rho_0 = 0.68$, $\rho_f = 0.68$, and $M = 13$.




Figure 5. Impulse Responses for the Behavioral Model with Different Price Patterns.

From the steady-state, two different sequences of shocks on A_t are imposed from $t = 2$ to $t = 6$, resulting in a steady rise in price (solid) and an increasing rise in price (dashed), respectively. We then plot, for these two cases, the net investment I_t and the output price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the behavioral model ($\alpha = 0.95$ and $\rho_A = 0.05$). The other parameter values are: $B = 0.02$, $\bar{A} = 12$, $\delta = 0.6\%$, $r = 0.5\%$, $k = 22.8$, $C = 100$, $P_r = 463$, $\sigma_\varepsilon = 4.25$, $p = 0.8$, $\rho_0 = 0.68$, $\rho_f = 0.68$, and $M = 13$.

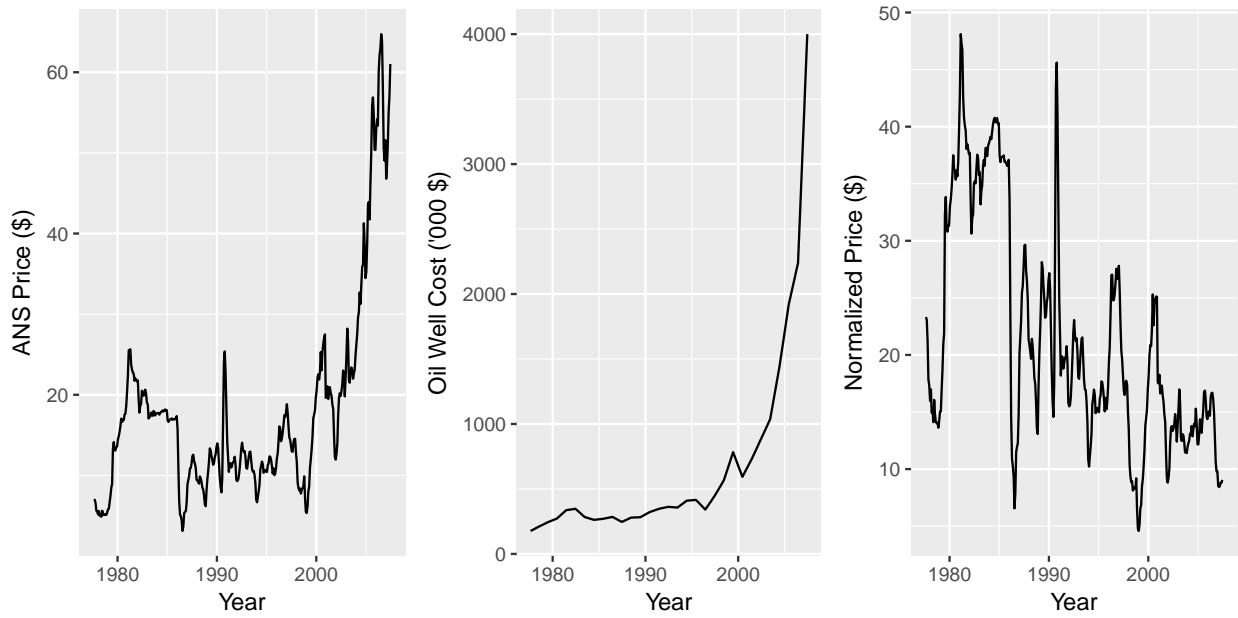


Figure 6. Alaska North Slope Oil Price and U.S. Well Cost Trend. The left panel of the figure plots monthly Alaska North Slope first purchase price; the middle panel plots the interpolated monthly U.S. well drilling cost (actual data is on an annual basis); the right panel plots the monthly price normalized by drilling cost (with the drilling cost of June 2000 as the basis). The sample period for all three plots is August, 1977 to June, 2007.


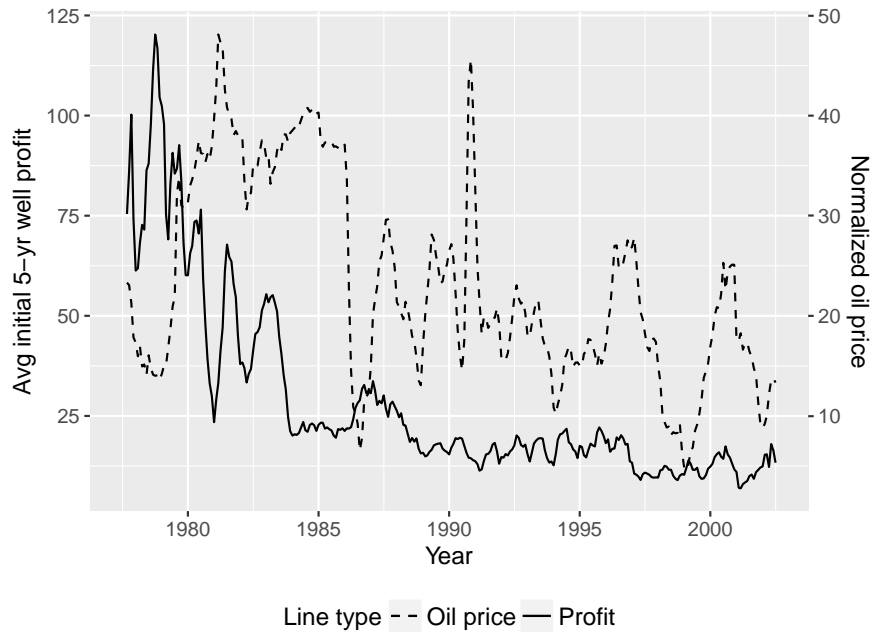
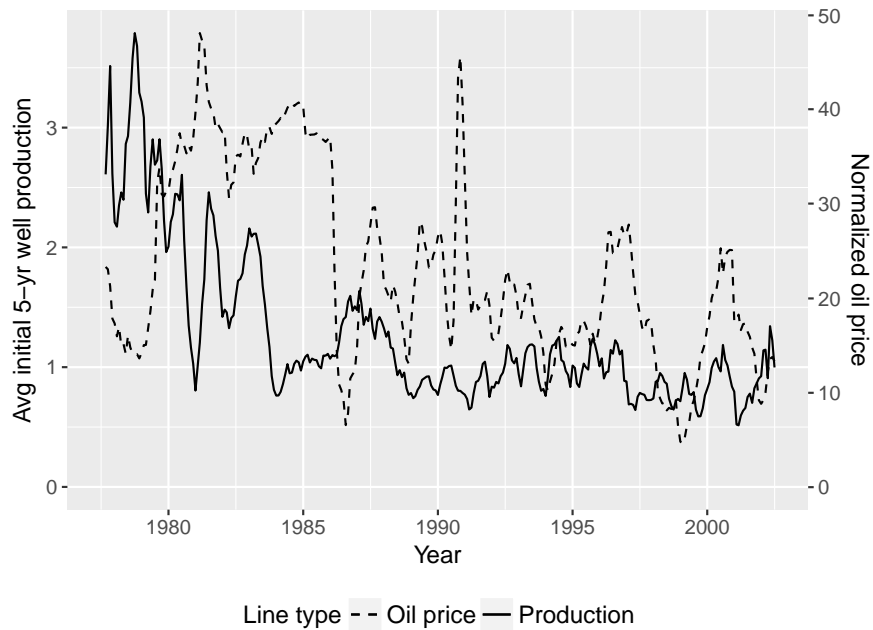


Figure 7. Well Investment Trend. The solid line shows the trend for the number of oil-related wells drilled by operators in Alaska in each month. The dashed line plots the time series of normalized oil price. The sample period is August, 1977 to June, 2007.



(a) Average initial five-year profit per well (\$MM)



(b) Average initial five-year production per well (MM bbl)

Figure 8. Well Production and Profit Trend. Figure 8(a) plots on a monthly basis the average profit (in million dollars) in the first 60 months of production (hence the initial 5 years) for wells that are drilled. Figure 8(b) plots on a monthly basis the average production (in million barrels) in the first 60 months of production for wells that are drilled. The dashed lines in both figures plot the time series of normalized oil price. The sample period for both plots is August, 1977 to June, 2002.

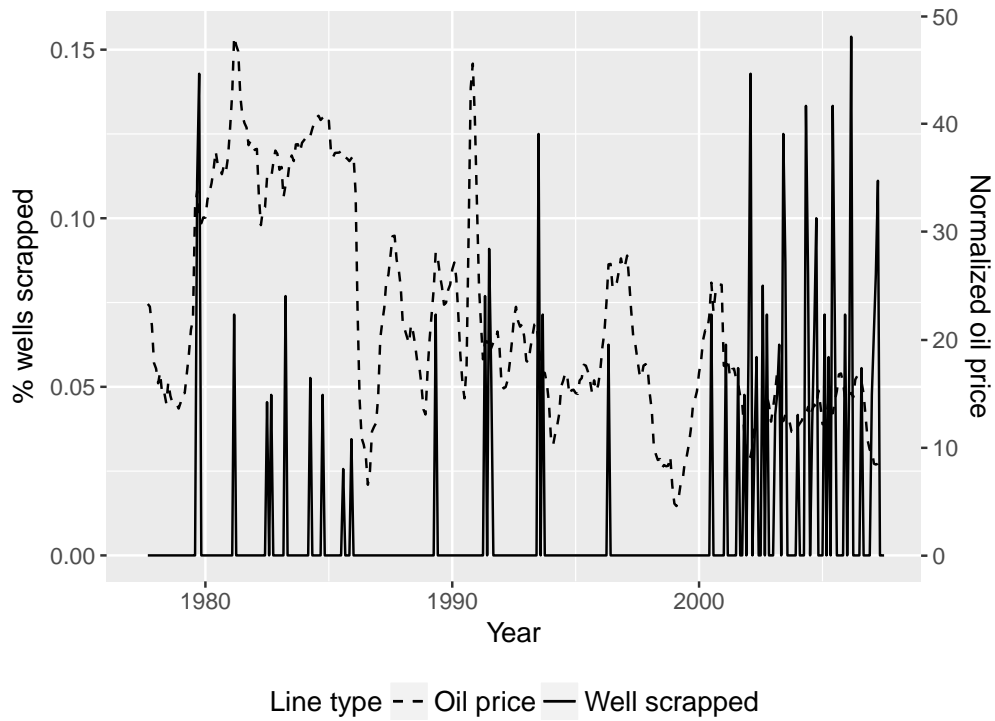


Figure 9. Well Scrappage Trend. The solid line plots the percentage of wells scrapped in each month among the wells that are given permits to drill. The dashed line plots the time series of normalized oil price. The sample period is August, 1977 to June, 2007.




Figure 10. Calibrated Impulse Responses for Historical Episode 1: the 2008 Great Recession. We plot the net investment I_t , total production Q_t , and the output price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the behavioral model ($\alpha = 0.9$ and $\rho_A = 0.25$). In the bottom panel, we also plot (on the right) the actual oil price from August 2008 to November 2011. From the steady-state, we choose a sequence of demand shocks on A_t so that the price pattern implied by the behavioral model roughly matches the actual oil price movements. The other parameter values are: $B = 0.02$, $\bar{A} = 12$, $\delta = 0.6\%$, $r = 0.5\%$, $k = 22.8$, $C = 100$, $P_r = 463$, $\sigma_\varepsilon = 4.25$, $p = 0.8$, $\rho_0 = 0.68$, $\rho_f = 0.68$, and $M = 13$.




Figure 11. Calibrated Impulse Responses for Historical Episode 2: the 1986 Saudi Oil Glut. We plot the net investment I_t , total production Q_t , and the output price H_t from $t = 1$ to $t = 15$ for both the benchmark model ($\alpha = 0$) and the behavioral model ($\alpha = 0.9$ and $\rho_A = 0.25$). In the bottom panel, we also plot (on the right) the actual oil price from February 1986 to September 1987. From the steady-state, we choose a sequence of demand shocks on A_t so that the price pattern implied by the behavioral model roughly matches the actual oil price movements. The other parameter values are: $B = 0.02$, $\bar{A} = 12$, $\delta = 0.6\%$, $r = 0.5\%$, $k = 22.8$, $C = 100$, $P_r = 463$, $\sigma_\varepsilon = 4.25$, $p = 0.8$, $\rho_0 = 0.68$, $\rho_f = 0.68$, and $M = 13$.

	<i>Dependent variable:</i>				
	Number of wells drilled				
	(1)	(2)	(3)	(4)	(5)
Normalized oil price					
1-month lag	0.070 (0.054)				
3-month lag		0.083 (0.055)			
6-month lag			0.109** (0.054)		
9-month lag				0.121** (0.052)	
12-month lag					0.127** (0.051)
Rate of change					
decreasing decline	-0.867 (1.022)	-1.241 (1.034)	-1.602 (1.108)	-1.384 (1.082)	-1.661 (1.065)
increasing decline	-1.094 (1.108)	-1.240 (1.110)	-0.434 (1.002)	-0.428 (0.965)	-0.927 (0.938)
decreasing rise	-1.914* (1.102)	-1.741 (1.149)	-1.675 (1.107)	-1.903* (1.056)	-2.174** (1.044)
increasing rise	-2.630*** (0.952)	-2.398** (0.943)	-2.174** (0.903)	-2.205*** (0.854)	-2.477*** (0.844)
Constant	13.528*** (1.290)	13.302*** (1.311)	12.690*** (1.303)	12.477*** (1.210)	12.573*** (1.121)
Observations	358	356	353	350	347
<i>R</i> -squared	0.063	0.073	0.093	0.108	0.123
Adjusted <i>R</i> -squared	0.050	0.059	0.079	0.095	0.110

Table I. Total Wells Permitted on Normalized Oil Prices and Price Trends. The table is based on a sample of 5,121 oil-related wells drilled in Alaska between August 1, 1977 and July 1, 2007. The lagged prices are the normalized oil prices 1, 3, 6, 9, and 12 months prior to the month of drilling. The rate of change categories are defined as periods with continuous increase or decrease for at least 3 periods and the rate of change follows a convex or concave shape. Numbers in the parentheses are Newey-West standard errors allowing for 6-month maximum lag in autocorrelation. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

	<i>Dependent variable:</i>				
	Initial five-year profit				
	(1)	(2)	(3)	(4)	(5)
Normalized oil price					
1-month lag	-0.106** (0.050)				
3-month lag		-0.108** (0.053)			
6-month lag			-0.258*** (0.083)		
9-month lag				-0.095 (0.110)	
12-month lag					-0.425** (0.192)
Rate of change					
decreasing decline	1.806 (1.164)	2.483** (1.019)	3.432*** (1.044)	1.857* (1.057)	2.240** (0.916)
increasing decline	1.350 (1.009)	2.910*** (1.033)	2.399*** (0.916)	1.965** (0.863)	3.004*** (0.993)
decreasing rise	3.151 (2.775)	2.900 (2.885)	2.494 (2.757)	3.116 (2.869)	3.209 (2.801)
increasing rise	1.995*** (0.572)	1.764*** (0.683)	1.527*** (0.560)	2.075*** (0.723)	2.142*** (0.682)
Constant	74.669*** (19.740)	60.102*** (14.882)	104.908*** (28.849)	116.458*** (27.434)	130.062*** (32.291)
Unit FE	Y	Y	Y	Y	Y
Drilling Year FE	Y	Y	Y	Y	Y
Observations	4,055	4,044	4,027	4,005	3,991
<i>R</i> -squared	0.368	0.368	0.361	0.364	0.357
Adjusted <i>R</i> -squared	0.345	0.345	0.338	0.342	0.335

Table II. Five-year Well Profit on Normalized Prices. This table is based on a sample of 4,055 oil-related wells with tract number drilled in Alaska between August 1, 1977 and July 1, 2002. Well profit is calculated by summing up the monthly production times monthly normalized price for each well in the first 60 months of production. The lagged prices are the normalized oil prices 1, 3, 6, 9, and 12 months prior to the month of drilling. The rate of change categories are defined as periods with continuous increase or decrease for at least 3 periods and the rate of change follows a convex or concave shape. Numbers in the parentheses are cluster-robust standard errors clustered by units. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

	<i>Dependent variable:</i>				
	Initial five-year production				
	(1)	(2)	(3)	(4)	(5)
Normalized oil price					
1-month lag	-0.004 (0.002)				
3-month lag		-0.004** (0.002)			
6-month lag			-0.009* (0.005)		
9-month lag				-0.002 (0.007)	
12-month lag					-0.010 (0.008)
Rate of change					
decreasing decline	0.096 (0.065)	0.119** (0.060)	0.153** (0.069)	0.100* (0.057)	0.109** (0.055)
increasing decline	0.079 (0.060)	0.129** (0.058)	0.104* (0.059)	0.092* (0.051)	0.116** (0.055)
decreasing rise	0.125 (0.138)	0.113 (0.142)	0.099 (0.140)	0.125 (0.142)	0.126 (0.139)
increasing rise	0.070** (0.034)	0.058 (0.039)	0.052 (0.036)	0.074* (0.041)	0.075** (0.037)
Constant	2.363*** (0.665)	1.900*** (0.491)	3.071*** (0.809)	3.329*** (0.737)	3.651*** (0.988)
Unit FE	Y	Y	Y	Y	Y
Drilling Year FE	Y	Y	Y	Y	Y
Observations	4,055	4,044	4,027	4,005	3,991
<i>R</i> -squared	0.258	0.257	0.251	0.250	0.244
Adjusted <i>R</i> -squared	0.231	0.229	0.224	0.224	0.218

Table III. Five-year Well Production on Normalized Prices. This table is based on a sample of 4,055 oil-related wells with tract number drilled in Alaska between August 1, 1977 and July 1, 2002. Well production is calculated by summing up the monthly production for each well in the first 60 months of production. The lagged prices are the normalized oil prices 1, 3, 6, 9, and 12 months prior to the month of drilling. The rate of change categories are defined as periods with continuous increase or decrease for at least 3 periods and the rate of change follows a convex or concave shape. Numbers in the parentheses are cluster-robust standard errors clustered by units. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

	<i>Dependent variable:</i>	
	Whether a well is scrapped	
	(1)	(2)
Normalized oil price	-0.0004* (0.0002)	-0.001*** (0.0001)
2-yr avg perc. change in price	-0.009*** (0.002)	-0.009*** (0.001)
Region = Other		0.0004 (0.020)
Region = Cook Inlet		0.988*** (0.019)
Constant	0.018*** (0.003)	0.022*** (0.004)
Operator FE	N	Y
Unit FE	N	Y
Observations	4,757	4,757
<i>R</i> -squared	0.002	0.272
Adjusted <i>R</i> -squared	0.001	0.248

Table IV. Likelihood to Scrap on Normalized Oil Price and Price Change after Approval. This table is based on a sample of 4,757 oil-related wells with tract number that are approved between August 1, 1977 and July 1, 2007. The two-year average percent change in price is calculated as the 2-year average normalized oil price after approval minus the normalized oil price at the time of approval, and then divided by the price at the time of approval. For region, the reference level is the North Slope and Beaufort Sea region. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Parameter	Value	Justification
\bar{A}	12 + Year fixed effect	Coefficients from regressing monthly price over productive well count
B	0.02	
ρ_0	0.68	Coefficient from regressing residual prices on its lag
σ_ε	4.25	Standard deviation of residual from the price residual regression
δ	0.6%	Inverse of the median length of life of development wells
p	0.8	Proportion of past wells producing at least 1000 barrels of oil
P_r	463	Coefficients from regressing well cost over new well count
k	22.8	
C	100	Umbrella term for all operating costs
M	13	Average of the median well production for each month

Table V. Calibrated Parameter Values. This is a list of parameters in the model that can be directly or indirectly inferred from the Alaska data. When point estimate of the parameters is available, we use the point estimate. When the estimated values fall within a range, we round to a reasonable number within the range. The operating costs here include lease operating expense, gathering, processing and transport expense, water disposal costs, and any general and administrative (G&A) costs.

	<i>Dependent variable:</i>	
	Alaska North Slope crude oil first purchase price	
	(1) OLS	(2) IV
Number of productive wells	-0.021*** (0.007)	-0.017** (0.007)
Constant	12.200*** (3.270)	11.400*** (3.620)
Year FE	Y	Y
Observations	465	464
<i>R</i> -squared	0.963	0.963
Adjusted <i>R</i> -squared	0.959	0.959

Table VI. Demand Regression. These regressions can be expressed as $H_t = \bar{A} - B \cdot Q_t + \epsilon_t$, where H_t is the per-barrel Alaska North Slope First Purchase price, and Q_t is the number of wells producing at least 2000 barrels in that month. The first column uses OLS regression. The second column uses IV regression where the instrument is the number of productive wells in the previous month.

	<i>Dependent variable:</i>	
	Residual from OLS	Residual from IV
	(1)	(2)
Lagged residual (from OLS or IV)	0.673*** (0.034)	0.684*** (0.034)
Constant	0.004 (0.198)	0.004 (0.196)
Observations	464	463
<i>R</i> -squared	0.452	0.467
Adjusted <i>R</i> -squared	0.451	0.466
σ_ε	4.25	4.22

Table VII. Residual Demand Regression.. The residuals are calculated as $H_t - \hat{A}_t - \hat{B}Q_t$ where $\hat{A}_t + \hat{B}Q_t$ comes from predictions from regressions in Table VI. Column 1 uses residuals from the OLS regression and column 2 uses residuals from the IV regression. In both columns, we regress the current period residual on the previous period's residual. σ_ε is the standard deviation of the residuals of each lagged residual regression.

	<i>Dependent variable:</i>
	Oil well drilling cost
Number of newly-drilled wells	11.399* (6.912)
Constant	462.581*** (104.387)
Observations	360
<i>R</i> -squared	0.008
Adjusted <i>R</i> -squared	0.005

Table VIII. Demand Regression Calibration. We regress the per-well drilling cost associated with oil wells on the number of newly-drilled wells in the current month. The intercept is the P_r and the coefficient associated with the number of wells is $k/2$. $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.