CEO Overconfidence and Innovation *

Alberto Galasso
University of Toronto

Timothy S. Simcoe
Boston University School of Management and NBER

November 14, 2009

Abstract

We develop a career concern model of innovation, where CEOs may innovate to signal their own ability to the market. In this model, CEOs who are excessively optimistic about their own ability are more likely to innovate. The effect is larger in more competitive industries. We test these predictions using a panel of large publicly traded firms from 1980 to 1994. Our measure of CEO overconfidence is the one proposed by Malmendier and Tate, and uses CEO's personal investments to capture their "revealed beliefs" about their own firms future performance. We find a robust positive association between overconfidence and citation-weighted patent counts, in both cross sectional and fixed-effect models. This effect is larger in more competitive industries. Our results suggest that overconfident CEOs are more likely to take their firms in a new technological direction.

Keywords: Innovation, R&D, CEO Overconfidence, Managerial Biases.

JEL Codes: D80, O31, O32, O33.

*We thank Ulrike Malmendier for providing the data that made this research possible and seminar participants at the University of Toronto for helpful suggestions. Addresses for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6, Canada; Boston University School of Management, 595 Commonwealth Avenue, Boston, MA 02215, USA. E-mail: alberto.galasso@rotman.utoronto.ca; tsimcoe@bu.edu.

1 Introduction

Individual decision makers deviate from economic rationality in systematic ways, such as time-inconsistency, loss aversion, framing and overconfidence (Rabin, 1998; DellaVigna, 2008). Economists have started to explore the practical effect of these psychological biases for industrial organization. For example, behavioral biases can influence pricing decisions (DellaVigna and Malmendier, 2006); entry decisions (Goldfarb and Xiao, 2009); labour productivity (Bandiera et al., 2005); bidding in auctions (Brown et al., 2009); union negotiations (Krueger and Mas, 2004); and corporate finance (Malmendier and Tate, 2005a, 2008).

Because of the great uncertainty and complexity associated with research and development, the innovation process provides a setting where we might expect key decision-maker's behaviors to deviate from the forecasts of traditional economic theory. In this paper we take a first step in exploring the interplay of behavioral biases and innovation by focusing on a specific psychological bias: overconfidence.

The standard economic model assumes that individuals are on average correct about the distribution of states of the word. However, many studies in applied psychology show that people tend to be over-confident when it comes to their own abilities. For example, a large proportion of individuals report having above the median driving skills (Svenson, 1981), a better than average ability to solve trivia quizzes (Moore and Cain, 2007) or a very good chance of getting the job they want (Weinstein, 1980).

The literature on overconfidence provides several indications that CEOs and other high-ranking executives should be particularly susceptible to this bias. For example, Camerer and Lovallo (1999) show that overconfidence is stronger among highly skilled individuals, while Moore and Kim (2003) show that the effect tends to be stronger when beliefs involve outcomes that are naturally complex and hard to compare. Malmendier and Tate (2008) show that CEO overconfidence has a significant impact on merger and acquisition activity, and Malmendier and Tate (2005a) find that investments by an overconfident CEO's are more sensitive to cash flows than is typical.

This paper examines the impact of CEO overconfidence on firm's innovation activities. We begin with a simple career concern model where innovation helps CEOs to signal their talent to the market. In our model, innovation is risky: if a research project is not successful, the market will perceive that the CEO has low talent, and the board may fire him. Because overconfident CEOs underestimate the likelihood of a failure, they innovate more than non overconfident CEOs. The model also predicts that the impact of overconfidence is stronger when product market competition is intense.

To test these predictions, we combine standard measures of innovation, based on US patent data, with a measure of CEO overconfidence developed in a series of papers by Malmendier and Tate (2005a, 2005b, 2008). The measure is constructed using CEOs personal investments to capture their "revealed beliefs" about their own firms future performance. Specifically, CEOs are classified as overconfident if they hold highly in-the-money stock options after they are fully vested. Our panel data regressions are based on a sample of 290 firms and 627 CEOs during the period 1980 to 1994. These are large firms, primarily from manufacturing and technology industries, where we observe significant patenting.

We find that the arrival of an overconfident CEO is correlated with a 28 percent increase in patent citations. This correlation is robust to controlling for firm fixed effects and for a number of time-varying firm and CEO characteristics. We find similar effects for patent counts, R&D investment and citations per issued patent. Moreover, we find that the correlation between CEO overconfidence and innovation is stronger when product market competition is more intense. The results hold both when we identify overconfident CEOs with a fixed effect and when their overconfidence is allowed to change over time.

1.1 Related Literature

This paper builds upon a rich empirical literature on corporate innovation. Pakes and Griliches (1980) were the first to use a patent production function to study the relationship between R&D expenditure and large-firm patenting. Their model was extended by and Hausman, Hall and Griliches (1984), and Blundell, Griffith and Van Reenen (1999). Patent production functions have been used to study the effects of competition (Aghion, Bloom, Blundell, Griffith and Howitt, 1995), R&D spillovers (Bloom, Schankerman and Van Reenen, 2009), institutional ownership (Aghion Van Reenen and Zingales, 2009) and the strengthening of intellectual property protection (Hall and Ziedonis, 2001).

We also contribute to an emerging literature at the intersection of industrial organization and behavioral economics (see Camerer and Malmendier (2007) for a survey). Several theories from this literature suggest that overconfidence can have positive effects. For example, in the model of Bernardo and Welch (2001) overconfident decision makers are less likely to imitate their peers, and more likely to explore their environment, which generates positive information externalities. Goel and Thakor (2008) develop a model where risk-averse CEOs under-invest, relative to their shareholders' optimum, while moderately overconfident CEOs increase firm value by pursuing risky projects. These papers suggest a provocative interpretation of our own results: where Malmandier and Tate (2005a, 2008) show that overconfident CEOs are likely to destroy value through unprofitable mergers and to engage in sub-optimal investment behavior, we find that the same individuals are also more entrepreneurial and innovative than their peers.

Finally, we note the existence of a parallel literature on overconfidence in entrepreneurship. For example, Astebro (2003) studies overconfidence in a sample of small Canadian inventors, and shows that many inventors pursue their ideas even after a professional review suggests they have very little commercial promise. Thompson and Klepper (2007) develop a model in which over-confident managers are more likely to form a spin-off. Our results suggest that similar biases produce measurable effects at the opposite end of the firm-size distribution.

The remainder of the paper is organized as follows. Section 2 presents the model and the empirical predictions. Section 3 describes the data, measures of overconfidence and our econometric framework. Section 4 describes the empirical results. Section 5 concludes.

2 A Model of Overconfidence and Innovation

Aghion, Van Reenen and Zingales (2009) extend the Holmstrom (1982) career concern model by allowing the manager to innovate in order to signal his ability. In this section, we develop a variant of their framework in which we introduce managerial overconfidence.

There are two periods, $t = 1, 2$. The firm is run by a CEO whose ability $\theta \in \left\{0, \overline{\theta}\right\}$ (where $\overline{\theta} > 0$) is unknown to the market, and to the CEO. The market's prior beliefs about CEO ability are:

$$\Pr_M(\theta{=}\overline{\theta}) = \Pr_M(\theta = 0) = \frac{1}{2}.$$

The CEO' beliefs about θ depend on whether they are over-confident. Specifically, we assume:

$$\Pr_C(\theta{=}\overline{\theta}) = \frac{1}{2}(1 + o)$$
$$\Pr_C(\theta{=}0) = \frac{1}{2}(1 - o)$$

where $0 \leq o \leq 1$ captures CEO overconfidence. When $o = 0$ the market and the CEO share prior beliefs, when $o > 0$ the CEO thinks that the market underestimates his expected talent.

In period 1, the CEO chooses whether or not to innovate. Let us indicate the innovation strategy by $i \in \{0, 1\}$. (One might think of this as a choice between taking the firm in a new direction, which leads a broad increase in exploration, versus sticking with an established strategy.) If the CEO does not innovate ($i = 0$) the revenue realization is equal to zero and no information is revealed about the CEO's ability. If the CEO does innovate ($i = 1$), he incurs an innovation cost, I, and the period 1 revenue realization is equal to:

$$y_1 = \begin{cases} 1 \text{ with probability } p \\ 0 \text{ with probability } 1-p \end{cases}$$

if $\theta = \overline{\theta}$ and it is equal to

$$y_1 = \begin{cases} 1 \text{ with probability } \alpha p \\ 0 \text{ with probability } 1-\alpha p \end{cases}$$

if the ability is low.

We define $\alpha \equiv 1 - \pi$ where π is a measure of product market competition so that the difference in ability is more pronounced when competition is intense.[1] The term α can be interpreted as a reduced form of an un-modeled race in which a patent is awarded to the best idea in a technology field. The greater the degree of competition, the lower the likelihood that CEOs with low ability will be able to come up with innovations that are superior to those of the competitors.[2]

Following Holmstrom (1982), we assume that the CEO operates in a fully competitive market, and that the second period income of the CEO is equal to the market perception of his expected ability, conditional on the information acquired in period one.

The timing of the game is as follows: (i) the CEO chooses whether to pay I and innovate; (ii) period 1 revenue is realized and observed by the market that updates its assessment of the CEO's talent; (iii) the CEO decides whether to leave the firm based on the comparison between his expected period 2 income and his outside option.

The outside option for a CEO is to reallocate to another sector. As in Aghion, Van Reenen and Zingales (2008), we assume that the ability is sector specific so that in the new sector the CEO receives a compensation independent of his current talent and equal to:

$$\underline{w} = \frac{1}{2}\overline{\theta} - \delta$$

where δ is the switching cost.

We solve the model by backward induction. If the CEO decides to innovate, market beliefs follow Bayes' rule. The CEO's income in period 2 if he remains in the firm is equal to:

---

[1] In the Appendix we show that qualitatively, results would be unchanged if one assumed that competition has an impact on high ability CEOs as long as this impact is not as large as the one on low ability CEOs.

[2] Consider this simple rent seeking game that Baye and Hoppe (2003) show to be strategic equivalent to the classic patent race model of Loury(1979). Two players H (high ability) and L (low ability) exert effort sustaining marginal costs c_H and c_L with $c_H < c_L$. The probability that each player obtains the patent is $x_i/(x_i + x_j)$. If there is no competition ($x_j = 0$) each player obtains the patent with probability 1. If the two players compete, they obtain the patent with probabilities $p_H = c_L/(c_L + c_H)$ and $p_L = c_H/(c_L + c_H)$. Because $p_L < p_H$ competition has a stronger impact on the low talent player than on the high talent player.

$$w_2(y_1) = \Pr\left(\theta = \overline{\theta} | y_1\right) \overline{\theta}.$$

This implies that:

$$w_2(y_1 = 1) = \frac{p}{p + \alpha p} \overline{\theta} = \frac{\overline{\theta}}{1 + \alpha} \tag{1}$$

and that

$$w_2(y_1 = 0) = \frac{1 - p}{2 - p - \alpha p} \overline{\theta}. \tag{2}$$

We make the following assumption:

$$\frac{\overline{\theta}}{1 + \alpha} > \frac{1}{2}\overline{\theta} - \delta > \frac{1 - p}{2 - p - \alpha p} \overline{\theta} \tag{A1}$$

which guarantees that the manager will leave if period 1 revenue equals to zero.[3]

In period 1, the CEO will innovate if his expected utility from innovation, $U(i = 1) - I$, exceeds the ex-ante utility from not innovating, $U(i = 0)$. Because without innovation the market does not update its beliefs we have that:

$$U(i = 0) = \frac{1}{2}\overline{\theta}.$$

If $i = 1$, the CEO's expected period 2 compensation is:

$$\begin{aligned} U(i = 1) &= \left[\frac{1}{2}(1 + o)p + \frac{1}{2}(1 - o)\alpha p\right] \frac{\overline{\theta}}{1 + \alpha} \\ &+ \left[\frac{1}{2}(1 + o)(1 - p) + \frac{1}{2}(1 - o)(1 - \alpha p)\right] \underline{w} \\ &= \frac{1}{2}p\overline{\theta} + \frac{1}{2}(2 - p - \alpha p)\underline{w} + o\frac{p}{2}(1 - \alpha)\left(\frac{\overline{\theta}}{1 + \alpha} - \underline{w}\right) \end{aligned} \tag{3}$$

where the first term of (3) is the ex-ante probability that the CEO assigns to a high revenue realization times $w_2(y_1 = 1)$ and the second term is the ex-ante probability of $y_1 = 0$ times CEO's outside option.

Because the CEO innovates when $U(i = 1) - I \geq U(i = 0)$ there will be innovation in equilibrium only if innovation costs are not too large, specifically:

[3] Because $\frac{1}{1+\alpha} > \frac{1}{2} > \frac{1-p}{2-p-\alpha p}$ there exists a non empty set of parameters $(\alpha, \overline{\theta}, \delta, p)$ which satisfy A1.

$$I \leq \widehat{I} = \frac{1}{2}p\overline{\theta} + \frac{1}{2}(2-p-\alpha p)\underline{w} - \frac{1}{2}\overline{\theta} + o\frac{p}{2}(1-\alpha)\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right). \tag{4}$$

Condition (4) yields two sets of testable implications. The first set of predictions relates to the direct effect of CEO overconfidence. Because

$$\frac{\partial \widehat{I}}{\partial o} = \frac{p}{2}(1-\alpha)\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right) > 0$$

innovation takes place for a larger range of innovation costs when the CEO is overconfident. We can write this result as:

Implication 1 *Overconfident CEOs are more likely to innovate than non-overconfident CEOs.*

Second, the model suggests an interplay of product market competition and innovation. The cross-partial derivative

$$\frac{\partial^2 \widehat{I}}{\partial o \partial \alpha} = -\frac{p}{2}\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right) - \frac{p}{2}\frac{\overline{\theta}(1-\alpha)}{(1+\alpha)^2} < 0$$

and the fact that $\alpha = 1 - \pi$ imply that overconfidence and competition are complements (i.e. $\partial^2 \widehat{I}/\partial o \partial \pi > 0$).

Implication 2 *The impact of CEO overconfidence is stronger when product market competition is higher.*

2.1 Discussion

The model builds on a number of assumptions which are worthy of additional discussion. First, we assumed that CEO talent is sector specific and that every time a CEO reallocates to another sector he experiences a new draw from the distribution of ability. In the Appendix we show that the model delivers the same set of testable implications if we assume that managerial ability is the same in all sectors. Intuitively, even if low revenue realization affects negatively both current and future compensations, overconfident CEOs will underestimate the likelihood of low revenue and therefore will be more likely to innovate. We show that the predictions hold even in the extreme case in which after a low revenue realization, the CEO leaves the firm and is never hired by other sectors i.e. $\underline{w} = 0$.

Second, we assumed that a CEO who reallocates to another sector sustains a switching cost equal to δ. In the Appendix we show that this assumption, despite reducing the likelihood of innovation, has no impact on the effect of overconfidence and on its interaction with product market competition. We also show that the two testable predictions are valid if the innovator does not sustain the private innovation cost I. In this case innovation occurs only if $U(i=1) \geq U(i=0)$ that is satisfied as long as the switching cost is below a threshold $\widehat{\delta}$. In the Appendix we show that this threshold increases in o and that $\partial^2\widehat{\delta}/\partial o \partial \pi > 0$.

Third, our baseline model assumed that there is no impact of competition on non-innovating CEOs. In the Appendix we extend the model assuming that in the absence of innovation the firm may experience a loss and that the likelihood of this loss is greater when competition is intense. In this case competition affects $U(i=0)$ because in the event of a loss the CEO has to reallocate to a different sector. We show that the two testable predictions of our baseline model hold in this alternative environment.

Finally, we assumed that high ability CEOs realize high revenue with probability p whereas low ability CEOs realize it with probability αp with α equal to $1-\pi$. In the Appendix we generalize the framework assuming that the probabilities of high revenue realization are $p(\pi)$ if $\theta = \overline{\theta}$ and $q(\pi)$ if $\theta = 0$ with $p(\pi)$ and $q(\pi)$ decreasing functions and $p(\pi) < q(\pi)$. We show that CEO overconfidence has a positive effect on innovation in this generalized setting. We also show that the effect of overconfidence increases with product market competition as long as product market competition has a stronger impact on low talent CEO than on high ability CEOs i.e. $q'(\pi) < p'(\pi) < 0$.

This generalized model highlights a distinction between our setting and the model of Aghion, Van Reenen and Zingales (2009) that assume $p = 1 - \pi$ (and that α does not depend on π). In their model talent is more valuable when competition is less intense (if $\pi = 1$ both types of CEOs realize zero revenue) whereas in our model talent is more valuable when competition is more intense. In the Appendix we show that in their setting, as in our model, innovation takes place for a larger set of innovation costs when the CEO is overconfident. Nevertheless, the impact of overconfidence on innovation is lower when π is large and the second testable prediction of our baseline model no longer holds. We also show that, despite the negative cross-partial derivative, once we fix the level of innovation cost we can find a threshold for product market competition, π^0, such that when $\pi \leq \pi^0$ both overconfident and non-overconfident CEOs innovate whereas if $\pi > \pi^0$ only overconfident CEOs innovate. This results implies that, as our baseline model, the model of Aghion, Van Reenen and Zingales (2009) is consistent with the idea that only overconfident CEOs innovate when competition is intense.

3 Data and Methods

3.1 Data

We begin with a panel of 450 large publicly traded U.S. firms between 1980 and 1994. These data are described in Hall and Liebman (1998) and Yermack (1995). Each firm in the sample appeared at least four times on a *Forbes* magazine list of the largest U.S. companies. These data provide a very detailed picture of CEO's stock option holdings, which Malmendier and Tate (2008) use to construct the measures of CEO overconfidence described below.

We use the Compustat firm identifier (GVKEY) to merge this panel of large publicly traded firms to the NBER US patent data file. The NBER patent data are described in Hall et al. (2001), and provide detailed information on all U.S. patents during our sample period, including application and grant years, citations to other patents, and assignee codes that can be used to identify ownership. To match U.S. patent assignee codes with Compustat firms, we started with the name-matching tool of Bessen (2009) and then searched by hand for variations on the firm names in our panel. After dropping firms in the Finance, Insurance and Real Estate sector (one-digit SIC code 5), which has a very low rate of patenting, we arrive at an estimation sample with 290 firms, 3,648 firm-years and 627 individual CEOs.[4] Table 1 provides summary statistics for this sample.

Our primary measure of innovation is a citation-weighted count of U.S. patents. This measure builds on a substantial literature that documents the link between patents, citations and firm value (Pakes and Griliches, 1980; Hall et al., 2005; Harhoff et al., 1999; Aghion et al., 2009, inter alia). Patents are assigned to a firm-year observation using their filing date, and we weight each patent by the truncation-adjusted citation count field contained in the NBER data (see Hall et al., 2001, for details).

We also consider several additional innovation metrics. First, we de-compose our primary measure into an unweighted patent count, and the average number of citations per patent (excluding self cites). Second, we use the research and development expenditures (Compustat item 46) as a measure of innovation inputs. Since firms are not required to account for their R&D expenditures, this variable has many missing values, even after we interpolate over any gaps of three years or less. Finally, in a series of extensions, we examine changes in originality and the share of self-citations. Table 1 shows that the distribution of innovative activity in our sample is highly skewed. While the median firm-year observation consists of a single patent that receives 6 citations, the sample mean is much higher, at 28 patents and 489 cites.

To measure competition, we use a Lerner index, as in Aghion et al. (2009). Specifically,

[4]Retaining firms from the FIRE sector does not change the main results.

we calculate the median gross margin of all firms in the Compustat database with the same two-digit SIC code as a focal firm. Our baseline model allows this competition measure to vary over time. However, we also consider robustness tests that use a time-invariant Lerner index, or a dummy for firms whose average gross margin over the entire sample period falls above the median of all firms the estimation sample.

As additional controls, we use on a variety of the accounting data reported by Compustat. Our main Compustat items are Sales (item 1); a capital-labor ratio constructed from the book value of total assets (item 6) and the number of employees (item 29); and a deflated R&D stock. To construct the R&D stock, we follow the method described in Hall (1990), depreciating all reported R&D activity at a rate of 15 percent over a ten year period. We also have several CEO-level control variables used in Malmendier and Tate (2008), including measures of stock and vested option holdings, age, job tenure, and a set of dummies categorizing their educational background as finance or technical. CEOs with a "finance" background received a degree in accounting, finance, business (including MBA) or economics. CEOs with a "technical" background received a degree in engineering, physics, chemistry, math, operations research, biology or applied sciences.

3.2 Measuring Overconfidence

Our measures of CEO overconfidence build on a series of papers by Malmendier and Tate. These papers use CEOs' personal investment decisions to construct a proxy for overconfidence, or systematic over-estimation of the returns to holding stock in their own firm. The key idea behind this measurement strategy is to focus on the decision to exercise executive stock options. These options give the holder a right to purchase stock in their own company, usually at the prevailing price on the date of the option grant. They typically have a ten year life, and are fully exercisable after a four year vesting period. At exercise, the shares are almost always immediately sold (Ofek and Yermack, 2000).

While investors may hold ordinary options because the right to delay a stock purchase has positive value (Merton, 1973), executive stock options have several unique features that create strong incentives for exercise, so long as they are fully vested (and in the money). In particular, executive stock options are non-tradable, and CEOs cannot legally hedge their risk by short-selling shares in their own firm. Moreover, most CEOs are highly exposed to idiosyncratic risk associated with their own firm through equity compensation, stock holdings and firm-specific human capital. Consequently, standard models of decision-making under uncertainty (e.g. Hall and Murphy, 2002) indicate that a risk-averse CEO should exercise vested executive options before expiration as long as the stock price is sufficiently high. Nevertheless, many of the

CEOs in our sample fail to exercise their executive options, often repeatedly. Malmendier and Tate use this behavior as an indicator of CEO overconfidence, or systematic over-estimation of expected returns from holding the stock.

While there are other potential explanations for CEO's decision to hold fully vested executive options, Malmendier and Tate (2008) provide strong evidence for the overconfidence interpretation. In particular, their research shows that failure to exercise in-the-money executive options is positively associated with value-destroying merger and acquisition activity, and a relatively high sensitivity of investments to cash flows. These findings are consistent with the idea that overconfident CEOs believe they can make good investments, but perceive the market price of debt financing as too high. Malmendier and Tate also find that CEOs do not earn abnormal returns from holding their executive options, relative to a benchmark case of exercising the options and investing the proceeds in an S&P 500 stock index. This suggests that "late" exercise does not reflect inside information about the future prospects of the company. After considering a variety of other interpretations (e.g. board pressure, risk-tolerance, taxes and procrastination) Malmendier and Tate (2008) argue that the broad pattern of results is most consistent with the idea that CEOs who fail to exercise their fully vested and in the money executive options are systematically over-estimating the future performance of their own firm, i.e. they are overconfident. Our paper builds on the measurement strategy of Malmendier and Tate (2008) to construct two proxies for CEO overconfidence:

Holder67 This indicator variable is identical to the Holder67 variable in Malmendier and Tate (2008). To construct this variable, they examine all CEO option packages five years before expiration (after they are fully vested). The variable Holder 67 equals one for any CEO that fails to exercise an executive option at that time after their stock price has risen by at least 67 percent. This 67 percent exercise threshold is calibrated using the Hall and Murphy (2002) framework, assuming that two-thirds of CEO wealth is tied to company stock. Under this framework, failing to exercise an option that is 67 percent in the money implies a constant relative risk-aversion parameter of three. This measure treats overconfidence as an absorbing state: a CEO can become overconfident (though switching from zero to one is very rare), but will never change back.

In our estimation sample, Holder67 classifies roughly half of all CEOs as overconfident. However, a large proportion of all CEOs are not classified, either because they served a short tenure (so there was no opportunity to exercise a fully vested option package), or because their stock price did not appreciate by 67 percent. Moreover, Holder67 is only defined for CEOs who have been with a company for at least four years. Thus, our estimation sample contains 1,344 observations where Holder67 is defined. One can think of Holder67 as identifying CEOs who become overconfident following a 67 percent increase in the stock price of their firm. Our second measure is motivated by the idea that overconfidence may be a permanent trait.

Overconfidence This measure is a CEO fixed effect that equals one for all CEOs where Holder67 equals one, and zero for all CEOs where Holder67 equals zero. In practical terms, Overconfidence is simply the maximum value of Holder67 for a given CEO. This is useful for models where we wish to exploit within-firm variation associated with the arrival of an overconfident CEO, as opposed to cross-sectional difference between firms. While Overconfident is defined for 2,230 observations in our sample, there are still 1,418 observations where it is undefined because of a short tenure or a stock that did not appreciate by at least 67 percent. Our main results are robust re-classifying these missing CEOs as non-onverconfident (though we have no justification for doing so).

Malmendier and Tate (2008) also use a variable called Longholder, which defines a CEO as overconfident if they hold an executive option until the year of expiration. We do not use this measure because many Longholder CEOs are in non-patenting sectors of the economy, so we are left with only 23 Longholder CEOs in our panel that actually receive a patent.

3.3 Methods

Our main econometric models focus on the relationship between count-based measures of innovative activity Y_{it} at firm i in period t, and measures of CEO overconfidence O_{it}. We typically model the conditional expectation of innovative activity as

$$E[Y_{it}] = \exp(\alpha O_{it} + x_{it-1}\beta + \gamma_i + \lambda_t) \tag{5}$$

where x_{it-1} is a vector of control variables (lagged one period to account for obvious forms of simultaneity), γ_i is a firm-specific idiosyncratic effect, and λ_t is a vector of time-period effects. Equation (5) uses the log-link formulation because of the non-negative and highly skewed nature of our count-based dependent variables. However, Wooldridge (1999) emphasizes that Poisson quasi maximum-likelihood estimation will yield consistent estimates as long as the conditional mean is correctly specified, making it equally appropriate for positive and continuously-valued variables, such as R&D. We allow for arbitrary heteroskedasticity and autocorrelation (i.e. clustering standard errors).[5]

When x includes measures of the firm's R&D stock, equation (5) can be interpreted as a "knowledge production function" that translates past research investments into new inventions. In that formulation, α indicates whether firms led by overconfident CEOs receive more cite-weighted patents per dollar of R&D expenditure, so it is a measure of efficiency. We also

[5]Our results are robust to clustering standard errors at two digits SIC codes (presented in tables), firms and CEO level.

estimate models that omit the R&D stock from x, in which case α measures the combined effect of changes in R&D stocks and inventive efficiency.

We consider two methods of introducing the firm-specific effects γ_i in equation (5). One approach relies on the fixed-effects Poisson estimator (Hausman et al., 1984), which is analogous to the familiar within-group OLS estimator, and assumes that the x_{it} are strictly exogenous. Since strict exogeneity is a strong assumption in our setting, our main method is the "mean scaling" estimator of Blundell et al. (1999), which provides consistent estimates under the weaker assumption of predetermined x_{it} (as long as the first-moments of the data are stable). This method uses pre-sample data on the dependent variable to construct a mean, which then enters the estimation directly (analogously to x_{it}) to account for initial conditions. Blundell et al. (1999) show that this approach performs well even with relatively short pre-sample periods. We use ten years of pre-sample data below.

Since our overconfidence measures are only defined for CEOs who hold executive options that are at least 67 percent in the money after five years, the main estimates are based on an unbalanced panel. In particular, a single firm may enter or exit the panel depending on whether we can measure the overconfidence of its CEO. In order to isolate the within-firm variation associated with "switching" to an overconfident CEO, we also conduct a robustness test that is meant to more closely resemble the familiar difference-in-differences estimator. To do this, we identify 28 cases where a company switches from a CEO that cannot be identified as overconfident to a CEO that we can identify as overconfident. Using equation (5), we compare the change in innovation following this switch to the change in innovation following a transition between two CEOs that cannot be identified as overconfident.

4 Results

4.1 Overconfidence and Innovation

Table 2 presents our first set of regression results, which show a robust positive association between CEO overconfidence and innovation. The dependent variable in all models is a cite-weighted patent count, or equivalently, a total citation count for the issued patents applied for in year t. All models in Table 2 are estimated via Poisson, with robust standard errors to account for over dispersion. Columns (1) through (5) use the *Overconfidence* measure, while models (6) and (7) consider the alternative Holder67, which leads to a smaller estimation sample.

We begin in column (1) with a pooled cross sectional model that includes only year and two-digit SIC code effects, along with the overconfidence measure. Exponentiating the coefficient of 0.67 suggests that the overconfident CEOs in our sample receive roughly twice as many cite-weighted patents as their non-overconfident counterparts.

In column (2) we introduce firm fixed effects using the mean scaling approach of Blundell et al. (1999). While the Overconfidence coefficient falls to 0.39, or a 48 percent difference in innovative output, the correlation between CEO overconfidence and citation-weighted patents remains quite strong. Column (3) adds controls for Sales, the firm's Capital to Labor ratio, the CEO's age, age squared, the CEO's tenure and tenure squared. This produces almost no change in the Overconfidence coefficient relative to the model containing only the pre-sample means of inventive output.

In columns (1) through (3), the Overconfidence coefficient α measures the joint impact of changes in efficiency (more output per dollar of R&D) and innovative intensity (greater spending on innovation). In column (4) we add the log of each firm's R&D stock, so the model becomes a patent production function, where α measures current patenting per dollar of lagged R&D spending. As expected, we observe a very robust positive correlation between past R&D and current patenting (see Hall et al., 2005). The coefficient on Overconfidence also declines by about 33 percent, to 0.246, indicating that Overconfident CEOs obtain 28 percent more cite-weighted patents per dollar of lagged R&D spending than their counterparts. This difference could reflect either a higher patent propensity among overconfident CEOs, or a change in the direction of innovative activity that leads to greater research productivity.

Column (5) re-examines the patent production function model using a fixed-effects Poisson estimator. Relative to the mean scaling estimator, we see a very large increase for the coefficient on sales, while the effect of R&D expenditures declines to essentially zero. Reassuringly, there is little change in the coefficient on overconfidence.

Finally, columns (6) and (7) estimate the same models as columns (3) and (4) using the alternative Holder67 measure of overconfidence. Since Holder67 is only defined starting in the year when a CEO holds a fully vested executive stock option that has appreciated by 67 percent or more, the sample size declines sharply. However, the pattern of results is very similar. While the coefficient on overconfidence is slightly greater, it still falls by about 20 percent when we move to a production function model that includes the R&D stock. Overall, the results in Table 2 document a strong positive association between overconfidence that is robust to a variety of measurement and empirical modeling strategies. We take these results as support of the first prediction in the theoretical model.

4.2 Alternative Measures

Table 3 asks whether our baseline results in Table 2 are driven by greater output (more patents), greater input (more R&D), or greater impact (more cites). We find the answer to be "yes" based on production function estimates with and without firm effects.

The first two columns in Table 3 use unweighted patent counts as the dependent variable. The results in column (1) suggest that overconfident CEOs file for about 16 percent more patents per dollar of lagged R&D stock than CEOs who are not overconfident. Adding fixed effects in column (2) causes the coefficients on sales and lagged R&D stock to fall, but does has no noticeable effect on the Overconfidence coefficient.

The middle two columns examine the link between Overconfidence and R&D. We drop the lagged R&D stock in this specification, since we are focused on inputs. In column (3), we find that overconfident CEOs perform about 19 percent more R&D than a typical CEO. Adding firm effects reduces this effect slightly (to 16 percent).

The last two columns in Table 3 examine the correlation between CEO Overconfidence and citations per patent. The results — both with and without firm effects — show a roughly twenty percent increase in the mean citation rate. Interestingly, there is little correlation between the firm level controls of Sales, the Capital-Labor ratio or the R&D stock and the average citation rate. We find these last results especially intriguing, as they evoke a change in innovative direction or impact, as opposed to merely an increase in the amount of R&D or patenting by overconfident CEOs.

4.3 Switching CEOs

While many of the panel regressions in Tables 2 and 3 include firm effects, they differ from a canonical difference-in-differences estimator in two ways. First, the panel excludes observations where our measure of CEO overconfidence is undefined. And second, while the mean scaling estimator allows us to include pre-determined (but not strictly exogenous) firm-level covariates, it does not isolate the within-firm co-variation of overconfidence and innovation (as evidenced by the fact that we can include the SIC effects). In Table 4, we use a sub-sample of the larger data set to conduct a within-firm analysis that identifies the Overconfidence effect purely from changes in innovative activity before and after CEO "switches."

For this analysis, we begin by identifying 28 cases where a CEO who was either not-overconfident or unclassified was replaced by an overconfident CEO (see Table B-1 for a list). In each case, we retained data for the four years preceding the switch and all years of data for the overconfident CEO. We performed a similar exercise for switches where a not-overconfident or unclassified CEO was replaced by a not-overconfident CEO. We use this dataset to conduct two types of analysis. First, we compare the change in innovation when the new CEO is overconfident to the change in innovation when the new CEO is not overconfident. Second, we consider the simple before/after comparison for switches that lead to an overconfident CEO.

In the first case, we include a separate time trend for the overconfident and control switches, to test for a difference in the innovation trends across firms prior to the arrival of a new CEO. All of these models are estimated in the fixed-effects Poisson specification to isolate within-firm variation, and we drop the firm-level controls which are unlikely to be strictly exogenous.

The first two columns in Table 4 present the results for citation-weighted patent counts. Column (1) shows the difference in differences results. Following a switch to an overconfident CEO, cite-weighted patent rise by 42 percent more than following a switch to a non-overconfident CEO. We cannot reject the hypothesis that there is no difference in the pre-switch patenting trends, although the estimated trend during that four year period is roughly twice as high for the firms that received an overconfident CEO. Column (2) focuses on the before versus after comparison within "treated" firms; a regression that would not be identified if all firms switched in the same year. This results show a large post-switch increase in cite-weighted patents.

Figure 1 provides an alternative look at the impact of a switch to an overconfident CEO on cite-weighted patents. Here, we allow the "treatment effect" to vary for each year, normalizing the coefficient for one year before the switch to zero. The figure shows that there is no discernible trend prior to the switch. In the year of the switch, there is a sharp increase, which doubles over the next two to four years, before leveling off.

Columns (3) and (4) in Table 4 examine unweighted patent counts. While we find evidence of an increase in patenting, it is less dramatic than the results for citation-weighted patents. The difference in difference estimates show that a 16 percent increase in patenting following a switch to an overconfident, CEO. This effect is not statistically different from zero. If we exclude the switches to a non-overconfident CEO, the point estimate on patenting increases to 26 percent and is significant at the 10 percent level.[6]

Finally, columns (5) and (6) examine changes in citations per patent. Here we find a large difference in the change between switches to overconfident and non-overconfident CEOs. In column (5), the coefficient on Overconfidence implies that the patent citation rate increases by 40 percent following the arrival of an overconfident CEO. The effect is small, though still significant at the 10 percent level when we focus on the before-after comparison.

4.4 Overconfidence and Competition

Table 5 presents several results related to the second prediction of our model. Specifically, the model suggests that the association between overconfidence and innovation will be stronger when firms face more competition. To examine this relationship, we interact the Overconfidence indicator variable with several variations on the Lerner index, or gross margin, which we assume

[6]Dropping the pre-switch trends in column (3) yields similar estimates.

is inversely related to product market competition. All of these regression use our baseline patent production function specification (see column (4) in Table 2), pre-sample mean scaling and cite-weighted patents as the outcome variable.

Column (1) uses a time-varying Lerner index calculated as the median gross margin of all firms in a particular two-digit SIC code. In this specification, the main effect of Overconfidence is economically large and statistically significant. While the main effect of the Lerner index is negative (less competition yields less innovation), the effect is not statistically significant. To provide a sense of the effect size, we note that a one standard deviation change in the Lerner index (or an additional 5 points of gross margin) is associated with a roughly 3 percent change in cite-weighted patents per dollar of R&D stock. However, the slope of this relationship between competition and innovation is roughly an order of magnitude larger for Overconfident CEOs. In particular, the interaction between Overconfidence and the Lerner index is large and statistically significant, as predicted by our model.

In column (2), we find a qualitatively similar pattern using the Holder67 measure of CEO overconfidence. The main effect of overconfidence is economically large and statistically significant. The main effect of competition is negligible. And the interaction is negative and significant. Once again, the interpretation is that the relationship between Overconfidence and citation-weighted patents is stronger for firms facing more competition.

Columns (3) and (4) return to our primary Overconfidence measure, but use different measures of competition. In column (3) we restrict the Lerner index to be constant over time, but continue to base the measure on the median gross margin of all Compustat firms in a two-digit SIC code. Note that we cannot estimate a main effect of competition in such a model, since the measure is collinear with industry effects. The results in column (3) are nevertheless very close to those in column (1).

Finally, column (4) measures competition using Lerner50, a dummy for firms with an average gross margin that is above the median gross-margin of all firms in our data set. Thus, we rely on within-sample variation in competition, rather than variation in the Compustat data measured at the industry-level. Once again, we find that the relationship between innovation and overconfidence is stronger when competition is more intense.

4.5 Extensions

Our final set of results use a series of alternative outcome variables to explore the idea that overconfident CEOs do not simply increase the level of innovation, but rather cause a change in the direction pursued by the firms they manage.[7] In columns (1) and (2), the outcome variable is an originality weighted patent count. Originality, as defined in Hall et al. (2001), is essentially one minus a Herfindahl of the concentration of a patent's backwards citations across classes. Thus, more original patents cite a more diverse array of prior art. The results in columns (1) and (2) show that originality weighted patent counts increase with CEO Overconfidence, and more so in industries with lower gross margins.

In columns (3) and (4), we use self-citations to construct a new measure called the Derivative Patent Share. We classify a patent as "derivative" if more than half of its total citation are to other patents assigned to the same firm, i.e. they are self cites. We then calculate the proportion of all patents that are "derivative" for a given fir-year and use that proportion as our outcome variable. Column (3) shows that there is no meaningful relationship between Overconfidence and the derivative patent share in the pooled panel regressions. However, when Overconfidence is interacted with competition, we find that derivative patenting declines for overconfident CEOs, but increase for overconfident CEOs when there is little competition. This result suggests that overconfident CEOs in profitable industries increase innovation, but focus on familiar problems. Overconfident CEOs in highly competitive fields appear to try for new innovations, perhaps in an effort to escape from the levels of competition at their current product-market location.

5 Conclusions

In this paper we study the relationship between CEO overconfidence and innovation. We use a simple career concerns model to show that CEO overconfidence can increase the incentive to innovate, and that its effect is stronger when product market competition is more intense. We find strong empirical support for the correlations predicted by the model. Overconfident CEOs are shown to systematically obtain more cited-weighted patents. Moreover we find that the effect of overconfidence increases with product market competition.

These strong correlations suggest that managerial behavioral biases may have a significant impact on the incentives to innovate at a firm level. They also suggest that applying tools from behavioral economics to innovation may yield novel insights into the determinants of firm R&D investments and patenting. We see our findings as complement to those in Aghion, Van Reenen and Zingales (2009). They show that institutional ownership, by reducing the likelihood of a CEO being dismissed after a decline in profits, encourages innovation. Our results suggest

[7] Table B-2 also presents a series of extensions that demonstrate the robustness of our main results to including extra covariates, changing the model specification or including dynamics.

that overconfidence, by reducing the ”perceived” likelihood of a CEO being dismissed, may encourage innovation.

References

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005, May). Competition and Innovation: An Inverted-U Relationship? *The Quarterly Journal of Economics 120*(2), 701–728.

Aghion, P., J. Van Reenen, and L. Zingales (2009). Innovation and Institutional Ownership. *NBER Working Paper 14769*.

Bandiera, O., I. Barankay, and I. Rasul (2005, August). Social Preferences and the Response to Incentives: Evidence from Personnel Data. *The Quarterly Journal of Economics 120*(3), 917–962.

Baye, M. R. and H. C. Hoppe (2003, August). The strategic equivalence of rent-seeking, innovation, and patent-race games. *Games and Economic Behavior 44*(2), 217–226.

Bernardo, A. E. and I. Welch (2001, 09). On the Evolution of Overconfidence and Entrepreneurs. *Journal of Economics & Management Strategy 10*(3), 301–330.

Bessen, J. (2009). Tool for matching assignee names. Technical report, NBER Patent Data Project.

Bloom, N., M. Schankerman, and J. Van Reenen (2005, February). Identifying Technology Spillovers and Product Market Rivalry. CEPR Discussion Papers 4912, C.E.P.R. Discussion Papers.

Blundell, R., R. Griffith, and J. van Reenen (1999, July). Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms. *Review of Economic Studies 66*(3), 529–54.

Brown, J., T. Hossein, and J. Morgan (2009). Shrouded Attributes and Information Suppression: Evidence from the Field. *The Quarterly Journal of Economics forthcoming*.

Camerer, C. and D. Lovallo (1999, March). Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review 89*(1), 306–318.

Camerer, C. and U. Malmendier (2007). *Behavioral Economics and its Applications*, Chapter Behavioral Economics of Organizations. Princeton University Press.

DellaVigna, S. (2008). Psychology and Economics: Evidence from The Field. *Journal of Economic Literature 47*(2), 315–372.

DellaVigna, S. and U. Malmendier (2006, June). Paying Not to Go to the Gym. *American Economic Review 96*(3), 694–719.

Goel, A. M. and A. V. Thakor (2008, December). Overconfidence, CEO Selection, and Corporate Governance. *Journal of Finance 63*(6), 2737–2784.

Goldfarb, A. and M. Xiao (2009). Who Thinks About the Competition? Managerial Ability and Strategic Entry in Local Telephone Markets. *University of Toronto Working Paper*.

Hall, B. (1990). The Manufacturing Sector Master File: 1959-1987. *NBER Working Paper No. 3366*.

Hall, B., A. Jaffe, and M. Trajtenberg (2005). Market Value and Patent Citations. *RAND Journal of Economics 36*, 16–38.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001, October). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. Working Paper 8498, National Bureau of Economic Research.

Hall, B. H. and R. H. Ziedonis (2001, Spring). The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995. *RAND Journal of Economics 32*(1), 101–28.

Hall, B. J. and J. B. Liebman (1998, August). Are CEOs Really Paid Like Bureaucrats? *The Quarterly Journal of Economics 113*(3), 653–691.

Hall, B. J. and K. J. Murphy (2002, February). Stock options for undiversified executives. *Journal of Accounting and Economics 33*(1), 3–42.

Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999). Citation frequency and the value of patented inventions. *Review of Economics and Statistics 81*(3), 511–515.

Hausman, J., B. Hall, and Z. Griliches (1984). Econometric Models for Count Data with an Application to the Patents-R&D Relationship. *Econometrica 52*(4), 909–938.

Holmstrom, B. (1999, January). Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies 66*(1), 169–82.

Klepper, S. and P. Thompson (2007). *Perspectives on Innovation*, Chapter Spinoff Entry in High-Tech Industries: Motives and Consequences (chapter 6), pp. 187–218. Cambridge University Press.

Krueger, A. B. and A. Mas (2004, April). Strikes, Scabs, and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires. *Journal of Political Economy 112*(2), 253–289.

Loury, G. C. (1979, August). Market structure and innovation. *The Quarterly Journal of Economics 93*(3), 395–410.

Malmendier, U. and G. Tate (2005a, December). CEO Overconfidence and Corporate Investment. *Journal of Finance 60*(6), 2661–2700.

Malmendier, U. and G. Tate (2005b). Does Overconfidence Affect Corporate Investment? CEO Overconfidence Measures Revisited. *European Financial Management 11*(5), 649–659.

Malmendier, U. and G. Tate (2008, July). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics 89*(1), 20–43.

Merton, R. C. (1973, Spring). Theory of Rational Option Pricing. *Bell Journal of Economics 4*(1), 141–183.

Moore, D. and T.-G. Kim (2003). Myopic Social Prediction and the Solo Comparison Paradox. *Journal of Personality and Social Psychology 85*, 1121–1135.

Moore, D. A. and D. M. Cain (2007, July). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes 103*(2), 197–213.

Ofek, E. and D. Yermack (2000, 06). Taking Stock: Equity-Based Compensation and the Evolution of Managerial Ownership. *Journal of Finance 55*(3), 1367–1384.

Pakes, A. and Z. Griliches (1980). Patents and R&D at the firm level: A first report. *Economics Letters 5*(4), 377–381.

Rabin, M. (1998, March). Psychology and Economics. *Journal of Economic Literature 36*(1), 11–46.

Svenson, O. (1981, February). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica 47*(2), 143–148.

Thomas, A. (2003). The Return to Independent Invention: Evidence of Risk Seeking, Extreme Optimism or Skewness-Loving. *The Economic Journal 113*, 226–239.

Weinstein, N. (1980). Unrealistic Optimism about Future Life Events. *Journal of Personality and Social Psychology 39*, 806–820.

Yermack, D. (1995). Do corporations award CEO stock options effectively? *Journal of Financial Economics 39*(2-3), 237–269.

Tables and Figures

Table 1: Summary Statistics

	Mean	Median	Min	Max	S.D.	Obs.
Weighted Patents	489.01	6.00	0.00	32509.00	1747.02	3648
log(Cites+1)	2.78	1.95	0.00	10.39	2.98	3648
Unweighted Patents	27.79	1.00	0.00	1221.00	81.29	3648
log(Patents+1)	1.51	0.69	0.00	7.11	1.80	3648
Cites per Patent	8.62	4.00	0.00	240.00	13.32	3648
log(R&D Expense)	3.80	3.92	0.00	8.73	1.94	1864
Overconfidence	0.58	1.00	0.00	1.00	0.49	2441
Holder67	0.49	0.00	0.00	1.00	0.50	1533
Lerner Index	0.11	0.09	0.03	0.22	0.05	3648
log(Sales)	7.85	7.75	2.95	11.81	1.12	3641
log(Employees)	2.68	2.72	-2.23	6.78	1.29	3627
log(Capital/Labor)	4.29	4.01	0.09	7.47	1.35	3637
log(R&D Stock)	2.47	0.00	0.00	10.36	3.00	3648
Construction	0.04					
Manufacturing	0.53					
Transport/Utilities	0.29					
Retail/Wholesale	0.11					
Services	0.03					
Pre-sample obs.	9.15					
Total Firms				290		
Total CEOs				627		
Overconfident				168		
Not-overconfident				136		
Unclassified				323		

Table 2: Overconfidence and Innovation

Poisson Panel Regressions

Unit of Observation = Firm-Year

Dependent Variable = Cite-weighted Patents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Overconfidence	0.674***	0.389***	0.360***	0.246**	0.241***		
	(0.19)	(0.12)	(0.14)	(0.11)	(0.07)		
Holder67						0.548***	0.411***
						(0.15)	(0.11)
ln(Sales)			0.414***	0.202**	0.892***	0.415***	0.056
			(0.12)	(0.10)	(0.17)	(0.13)	(0.14)
ln(Capital/Labor)			-0.062	0.088	-0.068	0.116	0.298**
			(0.18)	(0.14)	(0.11)	(0.23)	(0.15)
ln(R&D Stock)				0.324***	0.004		0.497***
				(0.09)	(0.06)		(0.07)
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SIC 2-digit Effects	Yes	Yes	Yes	Yes	N/A	Yes	Yes
CEO Controls	No	No	Yes	Yes	Yes	Yes	Yes
Firm Effects	No	BGV	BGV	BGV	Yes	BGV	BGV
Observations	2230	2230	2200	2200	1676	1344	1344
Firms	209	209	207	207	154	200	200
CEOs	279	279	277	277	211	270	270

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance. Holder67 is a dummy equal to 1 for all CEO years after the CEO fails to exercise an option 67% in the money with 5 years remaining duration. Overconfidence is the maximum value for Holder67 for a given CEO. BGV fixed effects are based on including pre-sample means of the dependent variable. CEO controls are Age, Age^2, Tenure and Tenure^2.

Table 3: Overconfidence and Innovation (Alternative Measures)

Poisson Panel Regressions

Unit of Observation = Firm-Year

Outcome Variable	Unweighted Patents		R&D Expense		Citations per Patent	
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.149*	0.154**	0.175*	0.151***	0.199***	0.202***
	(0.09)	(0.06)	(0.09)	(0.06)	(0.07)	(0.07)
ln(R&D Stock)	0.427***	0.277***			0.022	-0.000
	(0.06)	(0.10)			(0.10)	(0.10)
ln(Sales)	0.427***	0.160**	1.033***	0.773***	0.026	0.029
	(0.06)	(0.08)	(0.10)	(0.06)	(0.14)	(0.13)
ln(Capital/Labor)	0.079	0.101	-0.333*	-0.271***	0.038	0.049
	(0.11)	(0.15)	(0.19)	(0.09)	(0.11)	(0.10)
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes
SIC 2-digit Effects	Yes	Yes	Yes	N/A	Yes	Yes
CEO Controls	No	No	No	No	No	No
Firm Effects	No	BGV	No	Yes	No	BGV
Observations	2200	2200	1216	1199	2200	2200
Firms	207	207	123	119	207	207
CEOs	277	277	167	163	277	277

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance. Overconfidence is the maximum value for Holder67 for a given CEO where Holder67 is a dummy equal to 1 for all CEO years after the CEO fails to exercise an option 67% in the money with 5 years remaining duration. BGV fixed effects are based on including pre-sample means of the dependent variable.

Table 4: Overconfidence and Innovation (Diff in Diffs Analysis)

Poisson Panel Regressions

Unit of Observation = Firm-Year

Outcome Variable	Cite-Weighted Patents		Unweighted Patents		Cites per Patent	
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.354**	0.516***	0.164	0.257*	0.354***	0.247*
	(0.15)	(0.15)	(0.16)	(0.15)	(0.12)	(0.13)
Pre-Overconf x T	0.036		0.026		-0.093	
	(0.06)		(0.07)		(0.06)	
Pre-Switch x T	0.034		0.037		0.037	
	(0.04)		(0.03)		(0.04)	
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes
CEO Controls	No	No	No	No	No	No
Firm Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	560	280	560	280	560	280
Firms	52	28	52	28	52	28
CEOs	79	41	79	41	79	41

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance. Overconfidence is the maximum value for Holder67 for a given CEO where Holder67 is a dummy equal to 1 for all CEO years after the CEO fails to exercise an option 67% in the money with 5 years remaining duration. Pre-Switch is a dummy for the period prior to arrival of new CEO, and is interacted with a time-trend. Pre-verconf is a dummy for the pre-switch period for overconfident CEOs.

Table 5: Overconfidence Interactions with Competition

Poisson Panel Regressions

Unit of Observation = Firm-Year

Dependent Variable = Total Citations

	(1)	(2)	(3)	(4)
Overconfidence	0.744***		0.643***	0.330***
	(0.18)		(0.23)	(0.12)
Lerner Index	-0.487	-0.947		
	(4.12)	(4.07)		
Lerner x Overconfidence	-4.630***			
	(1.07)			
Holder67		0.690***		
		(0.19)		
Lerner x Holder67		-2.830**		
		(1.35)		
LernerSIC x Overconf			-3.598**	
			(1.49)	
Lerner50 x Overconf				-0.505**
				(0.24)
ln(R&D Stock)	0.337***	0.508***	0.338***	0.336***
	(0.08)	(0.08)	(0.08)	(0.08)
ln(Sales)	0.161	0.033	0.166	0.168
	(0.10)	(0.15)	(0.10)	(0.11)
ln(Capital/Labor)	0.073	0.294**	0.064	0.061
	(0.13)	(0.15)	(0.14)	(0.14)
Year Effects	Yes	Yes	Yes	Yes
SIC 2-digit Effects	Yes	Yes	Yes	Yes
CEO Controls	Yes	Yes	Yes	Yes
Firm Effects	BGV	BGV	BGV	BGV
Observations	2200	1344	2200	2200
Firms	207	200	207	207
CEOs	277	270	277	277

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance. Holder67 is a dummy equal to 1 for all CEO years after the CEO fails to exercise an option 67% in the money with 5 years remaining duration. Overconfidence is the maximum value for Holder67 for a given CEO. BGV fixed effects are based on including pre-sample means of the dependent variable. CEO controls are Age, Age^2, Tenure and $Tenure^2$. Lerner index measures of competition are described in the text.

Table 6: Overconfidence and Innovation: Extensions

Panel Regressions

Unit of Observation = Firm-Year

Outcome Variable	Originality Weighted Patents		Derivative Patent Share	
	Poisson (1)	Poisson (2)	OLS (3)	OLS (4)
Overconfidence	0.188**	0.406***	-0.006	-0.031**
	(0.08)	(0.13)	(0.01)	(0.01)
Lerner SIC x Overconf		-2.055*		0.275**
		(1.24)		(0.11)
ln(R&D Stock)	0.234***	0.241***	-0.004	-0.004
	(0.08)	(0.08)	(0.01)	(0.01)
ln(Sales)	0.187**	0.168**	0.008	0.009
	(0.08)	(0.09)	(0.01)	(0.01)
ln(Capital/Labor)	0.059	0.050	0.007	0.009
	(0.13)	(0.13)	(0.01)	(0.01)
Year Effects	Yes	Yes	Yes	Yes
SIC 2-digit Effects	Yes	Yes	N/A	N/A
CEO Controls	Yes	Yes	Yes	Yes
Firm Effects	BGV	BGV	Yes	Yes
Observations	2124	2124	1343	1343
Firms	199	199	155	155
CEOs	268	268	206	206

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance. Overconfidence is the maximum value for Holder67 for a given CEO where Holder67 is a dummy equal to 1 for all CEO years after the CEO fails to exercise an option 67% in the money with 5 years remaining duration. BGV fixed effects are based on including pre-sample means of the dependent variable. CEO controls are Age, Age^2, Tenure and Tenure^2. Lerner index measures of competition are described in the text. See Hall et al. (2001) for a definition of originality.

Figure 1: Annual Treatment Effects for Switches to Overconfident

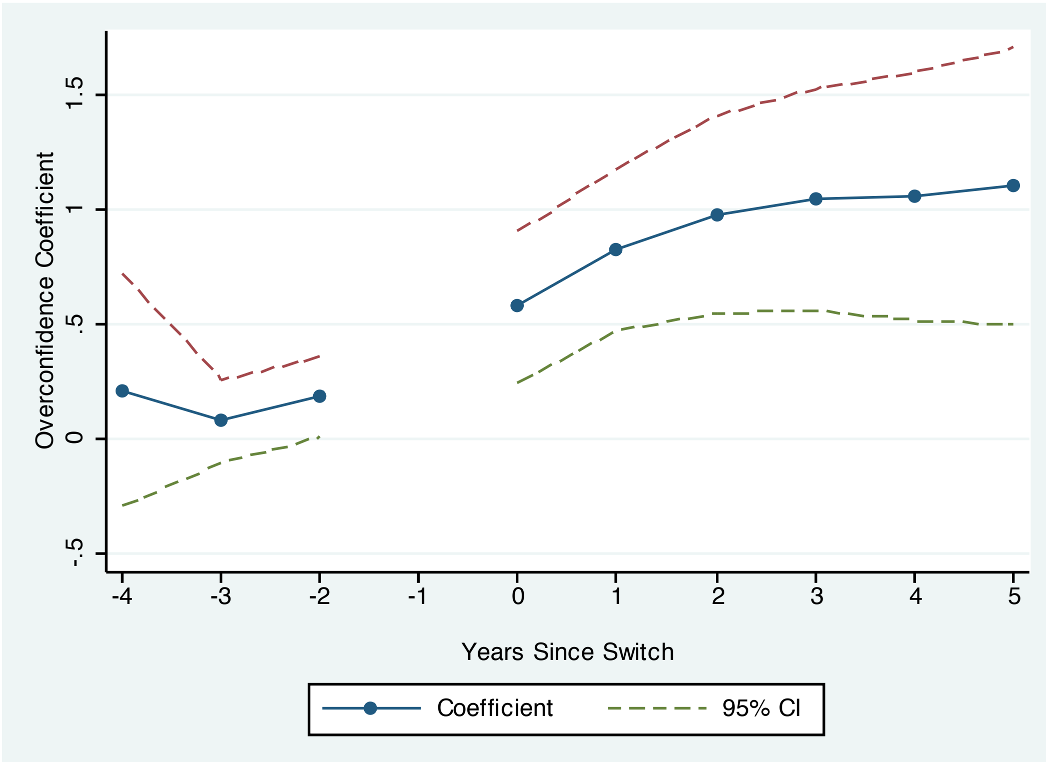

Appendix A: Extensions to the Theoretical Model

Non-Sector Specific Ability

In the baseline model we assumed that CEO talent is sector specific and that after relocating to a new sector the CEO receives a compensation that does not depend on his past performance:

$$\underline{w} = \frac{1}{2}\overline{\theta} - \delta.$$

We now relax this assumption and consider the case in which managerial ability is the same in all sectors. This implies that after low revenue realization the compensation obtained in a different sector will be

$$\underline{w} = w_2(y_1 = 0) - \delta = \frac{1-p}{2-p-\alpha p}\overline{\theta} - \delta$$

and that a CEO will never switch sector as long as $\delta > 0$. If $i = 1$, CEO's expected period 2 compensation is:

$$\begin{aligned} U(i=1) &= \left[\frac{1}{2}(1+o)p + \frac{1}{2}(1-o)\alpha p\right]\frac{\overline{\theta}}{1+\alpha} \\ &+ \left[\frac{1}{2}(1+o)(1-p) + \frac{1}{2}(1-o)(1-\alpha p)\right]\left(\frac{1-p}{2-p-\alpha p}\overline{\theta}\right) \\ &= \frac{1}{2}p\overline{\theta} + \frac{1}{2}(2-p-\alpha p)\frac{1-p}{2-p-\alpha p}\overline{\theta} \\ &+ o\frac{p}{2}(1-\alpha)\left(\overline{\theta}\frac{(1-\alpha)}{(\alpha+1)(2-p-\alpha p)}\right). \end{aligned} \tag{6}$$

Because innovation occurs only if $I \leq \widehat{I} = U(i=1) - U(i=0)$ and $U(i=0) = \overline{\theta}/2$ does not depend on overconfidence $\partial U(i=1)/\partial o \geq 0$ implies that $\partial\widehat{I}/\partial o > 0$. Moreover,

$$\frac{\partial^2 \widehat{I}}{\partial o \partial \pi} = p\overline{\theta}\frac{(1-\alpha)}{(\alpha+1)^2(p+p\alpha-2)^2}(3+\alpha-2p(1+\alpha)) \geq 0$$

so both of our testable implications hold in this alternative setting.

Aghion, Van Reenen and Zingales (2009) consider the extreme case in which, after a low revenue realization, the CEO leaves the firm and is never hired by other sectors i.e. $\underline{w} = 0$. In this case the ex ante compensation of a manager that innovates is:

$$U(i=1) = \frac{1}{2}p\overline{\theta} + o\frac{p}{2}(1-\alpha)\frac{\overline{\theta}}{1+\alpha}.$$

Notice that $U(i=1) - U(i=0) \leq 0$ if $o = 0$ (non-overconfident CEOs never innovate when

$\underline{w}$= 0) and that this difference becomes positive for overconfident CEOs as long as p is not too small. Moreover $\partial U(i = 1)/\partial o \geq 0$ and $\partial^2 U(i = 1)/\partial o \partial \pi \geq 0$ therefore this alternative model is also consistent with our testable predictions.

No Switching Cost

In the baseline model we assumed that when a CEO reallocates to another sector he sustains a switching cost δ. If CEOs can switch costlessy $\underline{w} = \overline{\theta}/2$. The absence of switching costs renders innovation more appealing because there is a higher payoff in the case of low revenue realization. Nevertheless, even in this alternative setting $\partial U(i = 1)/\partial o \geq 0$ and $\partial^2 U(i = 1)/\partial o \partial \pi \geq 0$. Therefore the assumption that $\delta > 0$ has no impact on our testable predictions.

No Innovation Cost

In the baseline model we assumed that the CEO sustains a private cost I when he innovates. If we remove this cost innovation occurs as long as $U(i = 1) \geq U(i = 0)$ that is satisfied as long as the switching cost is not too large:

$$\delta \leq \widehat{\delta} = p\theta\left(\alpha - 1\right)\frac{o + \alpha - o\alpha + 1}{2p - 4\alpha + 4p\alpha + 2p\alpha^2 + 2op - 2op\alpha^2 - 4}.$$

Also in this framework innovation takes place for a larger range of parameters when the CEO is overconfident:

$$\frac{\partial \widehat{\delta}}{\partial o} = p\theta \frac{\left(\alpha - 1\right)^2}{\left(\alpha + 1\right)\left(p + p\alpha + op - op\alpha - 2\right)^2}.$$

The cross partial derivative

$$\frac{\partial^2 \widehat{\delta}}{\partial o \partial \alpha} = p\theta \frac{\left(1 - \alpha\right)\left(2\alpha - 5p - 4p\alpha + p\alpha^2 - op + 2op\alpha - op\alpha^2 + 6\right)}{\left(\alpha + 1\right)^2\left(p + p\alpha + op - op\alpha - 2\right)^3} < 0$$

and therefore overconfidence and competition are complements (i.e. $\partial^2 \widehat{\delta}/\partial o \partial \pi > 0$) as in our baseline model. To see this notice that the second term in the denominator is negative because

$$\begin{aligned} p + p\alpha + op - op\alpha - 2 &= \\ p(1 + \alpha) + op(1 - \alpha) - 2 &\leq \\ 2p - 2 &\leq 0. \end{aligned}$$

The term $2\alpha - 5p - 4p\alpha + p\alpha^2 - op + 2op\alpha - op\alpha^2 + 6$ in the numerator is positive as long as

$$p \leq \widetilde{p}(o) = \frac{2\alpha + 6}{o(1-\alpha)^2 + 4\alpha - \alpha^2 + 5}$$

and because $\widetilde{p}'(o) < 0$ and $\widetilde{p}(1) = 1$ the term is positive for any value of α and o.

Competition Affects Non-Innovating Managers

We now relax the assumption that $U(i=0)$ is not affected by product market competition and extend the model assuming that competition affects non-innovating managers because it forces them to relocate to a different sector. We follow Aghion, Van Reenen and Zingales (2009) and assume that with probability $f(\pi)$ a non innovating firm incurs a loss and that $f'(\pi) > 0$. We also assume that that the CEO must relocate whenever the loss is incurred. In this case the CEO payoff without innovation is:

$$U(i=0) = (1 - f(\pi))\,\frac{\overline{\theta}}{2} + f(\pi)\underline{w}.$$

Because an increase in π increases the net gain $U(i=1) - I - U(i=0)$, competition renders innovation more appealing. It is important to notice that in this setting, as in our baseline model, $U(i=0)$ does not depend on o. This implies that the two testable predictions of our baseline model hold in this alternative environment.[8]

Generalization of the Competition Effect

In the baseline model we assumed that the difference in talent between high and low quality CEOs was captured by $\alpha = 1 - \pi$. We now generalize the framework by assuming that if the CEO innovates ($i = 1$), the period 1 revenue realization is equal to:

$$y_1 = \begin{cases} 1 \text{ with probability } p(\pi) \\ 0 \ \text{ with probability } 1 - p(\pi) \end{cases}$$

if $\theta = \overline{\theta}$ and it is equal to

[8] If the probability of incurring the loss is also affected by overconfidence (i.e. $f(o, \pi)$ with $\partial f/\partial o < 0$) an increase in o increases innovation as long as:

$$\left|\frac{\partial f}{\partial o}\right| < \frac{\frac{p}{2}(1-\alpha)\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right)}{\left(\frac{\overline{\theta}}{2} - \underline{w}\right)}.$$

$$y_1 = \begin{cases} 1 \text{ with probability } q(\pi) \\ 0 \text{ with probability } 1 - q(\pi) \end{cases}$$

if ability is low. We assume $q'(\pi) < p'(\pi) < 0$ and $q(\pi) \leq p(\pi)$: competition reduces the likelihood of high revenue but its effect is stronger for low ability CEOs. In this setting

$$w_2(y_1 = 1) = \frac{p(\pi)}{p(\pi) + q(\pi)}\overline{\theta} \tag{7}$$

and there is no change in

$$\underline{w} = \frac{1}{2}\overline{\theta} - \delta. \tag{8}$$

Notice that

$$U(i = 1) = \left[\frac{1}{2}(1 + o)p(\pi) + \frac{1}{2}(1 - o)q(\pi)\right] w_2(y_1 = 1) \tag{9}$$
$$+ \left[\frac{1}{2}(1 + o)(1 - p(\pi)) + \frac{1}{2}(1 - o)(1 - q(\pi))\right] \underline{w}$$

and that the CEO innovates when $U(i = 1) - I \geq U(i = 0)$. Therefore innovation occurs in equilibrium only if innovation costs are not too large, specifically:

$$I \leq \widehat{I} = U(i = 1) - U(i = 0).$$

Because

$$\frac{\partial \widehat{I}}{\partial o} = \frac{1}{2}\left(p(\pi) - q(\pi)\right)\left(w_2(y_1 = 1) - \underline{w}\right) > 0$$

the first testable prediction holds in this generalized setting. Moreover:

$$\frac{\partial^2 \widehat{I}}{\partial o \partial \pi} = \frac{1}{2}\left(w_2(y_1 = 1) - \underline{w} + (p(\pi) - q(\pi))\frac{p(\pi)}{(p(\pi) + q(\pi))^2}\right)\left[p'(\pi) - q'(\pi)\right] > 0$$

because $p'(\pi) - q'(\pi) > 0$ so the second testable prediction is also valid.

Competition reduces the impact of talent

In the baseline model we assumed that $\alpha = 1 - \pi$. This assumption implies that product market competition affects the probability that low ability CEOs have of generating high revenue from an innovation and that talent is more valuable in a competitive environment rather than in a non-competitive environment. We now follow Aghion, Van Reenen and Zingales (2009) and assume that $p = 1 - \pi$ and that α does not depend on π. Notice that in this variant of the model talent is more valuable when competition is less intense.

In this alternative setting the direct effect of overconfidence on innovation is analogous to the one in our baseline model:

$$\frac{\partial \widehat{I}}{\partial o} = \frac{p}{2}(1-\alpha)\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right) > 0.$$

Nevertheless, because talent is less valuable when competition is intense, the cross partial derivative has the opposite sign than the one in our model:

$$\frac{\partial^2 \widehat{I}}{\partial o \partial \pi} = -\frac{1}{2}(1-\alpha)\left(\frac{\overline{\theta}}{1+\alpha} - \underline{w}\right) < 0.$$

Therefore the impact of overconfidence on innovation is lower when π is large and the second testable prediction of our baseline model no longer holds.

It is important to notice that, despite the negative cross partial derivative, this alternative framework can still provide support to the idea that only overconfident CEOs innovate when competition is intense. To see this let us fix the level of innovation cost I. Notice that $U(i = 1) - I > U(i = 0)$ as long as

$$p > \widehat{p}(o) = \frac{(2I - 2\underline{w} + \overline{\theta})}{(1+\alpha+o(1-\alpha))} \frac{(1+\alpha)}{\left(\overline{\theta} - \underline{w}(1+\alpha)\right)}$$

and that $\widehat{p}$ is decreasing in o. This implies that non overconfident CEOs will innovate only if the level of product market competition is below $1 - \widehat{p}(0) \equiv \pi^0$. Therefore when $\pi \leq \pi^0$ both overconfident and non-overconfident CEOs innovate whereas if $\pi > \pi^0$ only overconfident CEOs innovate.

Appendix B: Additional Empirical Results

Table B-1: Overconfident CEOs in the Diff in Diffs Analysis

	Company Name	CEO Name	Year Hired
1	ALCOA	O'Neill Paul H.	1987
2	AMAX INC	Born Allen	1986
3	AMERICAN BRANDS INC	Alley William J.	1987
4	AT&T CO	Allen Robert E.	1988
5	C B I INDUSTRIES INC	Jones John E.	1989
6	CABOT CORP	Bodman Samuel W.	1988
7	CATERPILLAR INC	Fites Donald V.	1990
8	COLGATE PALMOLIVE CO	Mark Reuben	1984
9	COORS ADOLPH CO	Coors Jeffrey H.	1985
10	DEERE & CO	Bechere Hans W.	1990
11	E G & G INC	Kucharski John M.	1987
12	ENGELHARD CORP	Smith Orin R.	1984
13	GOODYEAR TIRE & RUBR CO	Mercer Robert E.	1983
14	GULF & WESTERN INDS INC	Davis Martin S.	1983
15	HERSHEY FOODS CORP	Zimmerman Richard A.	1984
16	INTEL CORP	Grove Andrew S.	1987
17	INTERNATIONAL PAPER CO	Georges John A.	1985
18	MEAD CORP	Roberts Burnell R.	1981
19	MOTOROLA INC	Fisher George M. C.	1988
20	NICOR INC	Cline Richard G.	1986
21	P P G INDUSTRIES INC	Sarni Vincent A.	1985
22	PACIFIC TELESIS GROUP	Ginn Sam L.	1988
23	RORER GROUP INC	Cawthorn Robert E.	1985
24	RYDER SYSTEMS INC	Burns M. Anthony	1983
25	SEARS ROEBUCK & CO	Brennan Edward A.	1986
26	TAMPAX INC	Shutt Edwin H.	1982
27	TOSCO CORP	O'Malley Thomas D.	1990
28	U S WEST INC	McCormick Richard D.	1991

Table B-2: Overconfidence and Innovation: Robustness Checks

Panel Regressions

Unit of Observation = Firm-Year

Outcome Variable	Cites	Cites	Cites	ln(Cites+1)	Cites
Specification	Poisson	Poisson	Poisson	OLS	NegBin
	(1)	(2)	(3)	(4)	(5)
Overconfidence	0.241**	0.251**	0.063*	0.313***	0.509***
	(0.12)	(0.11)	(0.03)	(0.09)	(0.17)
ln(R&D Stock)	0.323***	0.383***	0.129***	0.589***	-0.145
	(0.09)	(0.08)	(0.03)	(0.05)	(0.09)
ln(Sales)	0.201**	0.157	-0.002	0.257***	0.431***
	(0.10)	(0.15)	(0.04)	(0.07)	(0.10)
ln(Capital/Labor)	0.089	0.121	0.071	-0.052	-0.097
	(0.14)	(0.11)	(0.06)	(0.07)	(0.12)
Vested Options	0.004***				0.004***
	(0.00)				(0.00)
Stock Ownership	-0.626				-2.790*
	(2.30)				(1.52)
Finance Ed.		-0.072			
		(0.19)			
Technical Ed.		0.143**			
		(0.06)			
Chairman and Pres.		0.042			
		(0.06)			
ln(Cites t-1)			0.812***		
			(0.03)		
Year Effects	Yes	Yes	Yes	Yes	Yes
SIC 2-digit Effects	Yes	Yes	Yes	Yes	Yes
CEO Controls	Yes	Yes	Yes	Yes	Yes
Firm Effects	BGV	BGV	BGV	No	BGV
Observations	2200	1491	2037	1338	2411
Firms	207	153	206	155	227
CEOs	277	177	275	206	301

Cluster robust standard errors in parentheses: *10% significance; **5% significance; ***1% significance.