

# Overconfidence in Labor Markets\*

Luís Santos-Pinto<sup>†</sup>      Leonidas Enrique de la Rosa<sup>‡</sup>

February 3, 2020

## Abstract

This chapter reviews how worker overconfidence affects labor markets. Evidence from psychology and economics shows that in many situations, most people tend to overestimate their absolute skills, overplace themselves relative to others, and overestimate the precision of their knowledge. The chapter starts by reviewing evidence for overconfidence and for how overconfidence affects economic choices. Next, it reviews economic explanations for overconfidence. After that, it discusses research on the impact of worker overconfidence on labor markets where wages are determined by bargaining between workers and firms. Here, three key questions are addressed. First, how does worker overconfidence affect effort provision for a fixed compensation scheme? Second, how should firms design compensation schemes when workers are overconfident? In particular, will a compensation scheme offered to an overconfident worker have higher- or lower-powered incentives than that offered to a worker with accurate self-perception? Third, can worker overconfidence lead to a Pareto improvement? The chapter continues by reviewing research on the impact of worker overconfidence on labor markets where workers can move between firms and where neither firms nor workers have discretion over wage setting. The chapter concludes with a summary of its main findings and a discussion of avenues for future research.

## 1 Introduction

This chapter reviews how worker overconfidence affects labor markets. Workers' beliefs about their skills are at the heart of many labor market decisions. For example, the choice between a compensation scheme based

---

\*This working paper is a preliminary version of our chapter, with the same title, to appear in the *Handbook of Labor, Human Resources and Population Economics*, published by Springer-Verlag. We would like to thank Adrian Bruhin, José Mata, Joel Sobel, the editor, Marie Claire Villeval, and an anonymous referee for helpful comments and suggestions along the way.

<sup>†</sup>Corresponding author. Faculty of Business and Economics (HEC Lausanne), University of Lausanne, Bâtiment Internef 535, CH-1015, Lausanne, Switzerland. Ph: +41216923658. E-mail address: LuisPedro.SantosPinto@unil.ch

<sup>‡</sup>Department of Economics and Business Economics, Aarhus BSS, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark. Ph: +45 8716 5523. E-mail address: delaRosa@econ.au.dk

on absolute (e.g., a piece rate) or relative performance (e.g., a tournament) depends on beliefs about absolute and relative skill. The decision of how much effort to exert in jobs where performance depends on effort and ability also depends on beliefs about skill. Occupational and career choices also depend on beliefs about skill.

A large body of research in psychology and economics shows that in many situations, individuals are overconfident, that is, they overestimate their skills, performance, or desirable personality traits. Overconfidence is among the most well-documented judgement biases in the psychology literature. According to Myers (1998, p. 440): “on nearly any dimension that is both subjective and socially desirable, most people see themselves as better than average” and according to De Bondt and Thaler (1995, p. 389): “Perhaps the most robust finding in the psychology of judgment is that people are overconfident.” There are three main forms of overconfidence which are labeled by Moore and Healy (2008) as *overestimation*, *overplacement*, and *overprecision*.

Overestimation is the tendency to overestimate one’s absolute skills, performance, or desirable personality traits. Most people tend to have an overly positive self-image of their desirable personality traits (Taylor and Brown, 1988). Overestimation is also present in the workplace. Chief executive officers (CEOs) of large U.S. companies generally overestimate their ability to raise the company’s stock price and consequently they hold their stock options for too long (Malmendier and Tate, 2005, 2008; Malmendier et al., 2011). Truck drivers systematically and persistently overestimate the number of miles they drive each week (Hoffman and Burks, 2019).

Overplacement is the tendency to overestimate one’s relative skills, performance, or desirable personality traits. It is also referred to as the “better-than-average” or “Lake Wobegon” effect. The majority of managers (Larwood and Whittaker, 1977) and college professors (Cross, 1977) state they do above-average work. White-collar clerical and technical workers from the General Electric Company (Baker et al., 1988) and engineers in high-tech companies (Zenger, 1992) overestimate relative performance at their jobs. In Australia, the majority of people rate their job performance as above average (Myers, 1998). In New Orleans, a majority of federal magistrate judges place themselves in the lowest quartile according to the rate at which their decisions are overturned during their careers (Guthrie et al., 2001).

Overprecision is the tendency to overestimate the precision of one’s estimates or knowledge. It is also referred to as miscalibration and can be seen as a tendency to overestimate one’s forecasting ability or one’s knowledge. Most people provide overly narrow confidence intervals around numerical forecasts and display too low percentages of correctly answered items for a given confidence category (Lichtenstein et al., 1982; Hoffrage, 2016). People are more confident of their predictions in fields where they have self-declared expertise (Heath and Tversky, 1991): clinical psychologists overestimate the chances their predictions are accurate (Oskamp, 1965), physicians’ diagnoses of pneumonia are made with 88 percent confidence but only turn out to be right 20 percent of the time (Christensen-Szalanski and Bushyhead, 1981), and top U.S. corporate finance officers’

(CFOs) estimates of 80 percent confidence intervals for the expected annual S&P 500 returns only contain the actual realized returns 33 percent of the time (Ben-David et al., 2013).

Overconfidence has consequences for economic behavior which have been detected both in the laboratory and in the field. In the laboratory, participants in market entry experiments who overplace themselves are more likely to self-select into markets where rewards depend on relative performance (Camerer and Lovo, 1999). Participants in experimental asset markets who overestimate the precision of their information are more prone to display the winner's curse and to have lower trading performance (Biais et al., 2005) and tend to trade excessively (Deaves et al., 2009). Participants who overplace themselves are more likely to select being paid according to a tournament than either by a piece rate (Niederle and Vesterlund, 2007) or by a fixed payment independent of output (Dohmen and Falk, 2011). Overestimation and overprecision affect innovative decisions (Herz et al., 2014). Over- and underplacement influence individual risk taking (Bruhin et al., 2018).

Overconfidence also has consequences for economic behavior in the field. Consumers who overestimate their self-control overpay for gym memberships that they underutilize (DellaVigna and Malmendier, 2006) and those who overestimate the precision of their cell phone demand prediction pay more for their service (Grubb, 2009). CEO overconfidence, measured as the persistent failure to reduce personal exposure to firm specific risk, is a time-persistent personal characteristic and overconfident CEOs are 65 percent more likely to complete mergers, overpay for the target companies, and undertake value-destroying mergers (Malmendier and Tate, 2005, 2008). CFO overprecision is correlated with own-firm project overprecision and increased corporate investment (Ben-David et al., 2013). Truck drivers who overestimate the number of miles they drive each week are less likely to resign which benefits firms by increasing the profitability of training (Hoffman and Burks, 2019).

This chapter focuses on how worker overconfidence affects workers' decisions, firm behavior, and labor market outcomes. Regarding workers' decisions, the chapter considers the consequences of worker overconfidence for effort provision within a job, choice of compensation scheme, effort spent in searching for a new job, and occupational and career choices. Regarding firm behavior, the chapter discusses the impact of worker overconfidence on the type of compensation schemes offered to workers, which workers are promoted, how different workers are assigned to different jobs, and which workers are hired. In addition, the chapter describes how worker overconfidence influences labor market outcomes such as firm profit, worker utility, and social welfare.

The chapter is structured as follows. Section 2 surveys empirical research that documents overconfidence in general and highlights the differences between psychology and economic research on overconfidence. Section 3 reviews the economic explanations for overconfidence. Section 4 discusses theoretical and empirical research on the impact of worker overconfidence on internal labor markets, that is, labor markets where wages are determined by bargaining between workers and firms. Section 5 moves on to theoretical and empirical research on the impact of worker

overconfidence on external labor markets, that is, competitive labor markets where workers can move between firms and where neither firms nor workers have discretion over wage setting. Section 6 summarizes the main findings of the chapter and discusses avenues for future research.

## **2 Overconfidence: Evidence**

This section reviews empirical research on overconfidence.

### **2.1 Evidence from Psychology**

Psychology research on overconfidence focuses on three objectives. The first is documenting overconfidence. The second is identifying variables that moderate it. The third is proposing and testing alternative explanations for why it occurs. This review concentrates mainly on the first two objectives. Psychology explanations for overconfidence are briefly discussed in Section 3.

#### **2.1.1 Overestimation**

To measure overestimation psychologists ask participants to take a skill test or perform a task and then compare the outcome to participants' self-beliefs. For example, if a participant took a 20-item quiz, believes to have answered 15 items correctly, but only got 10 correct, she overestimated her absolute performance on the quiz by 5 questions.

The main findings of psychology research on overestimation are as follows. First, in many instances individuals overestimate their absolute skills but there are also many instances where they underestimate them (Moore and Healy, 2008). Second, overestimation at the individual level decreases with an increase in objective ability (Preston and Harris, 1965; Kruger and Dunning, 1999; Haun et al., 2000; Hodges et al., 2001). Third, individuals tend to underestimate their absolute performance on easy tasks and overestimate it on hard tasks (Moore and Cain, 2007; Moore and Healy, 2008); as Erev et al. (1994) point out, if people are uncertain about their absolute abilities, then there is more opportunity to underestimate (overestimate) one's absolute ability when performance is high (low), which we expect in easy (hard) tasks.

#### **2.1.2 Overplacement**

Psychologists have also been interested in documenting overplacement. Overplacement, like overestimation, can be measured at the individual and at the group levels. To measure overplacement at the individual level, each participant is asked to assess her relative skill and take a skill test, which provides an external and objective measure of her relative skill. A participant overestimates her relative skill when her relative performance falls short of her subjective assessment. To measure overplacement at the group level, participants' subjective assessments are converted into a statistic which is then compared to an aggregate objective criterion. For example, Svenson (1981) asked a sample of American drivers to rank their

driving skill and safety into one of ten deciles. The experiment documents overplacement at the group level: 46.3 percent placed themselves in the top two deciles on skill, and 82.5 percent placed themselves in the top three deciles in terms of safety.

The main findings of psychology research on overplacement are as follows. First, individuals overplace themselves on the vast majority of skills and desirable personality traits (Taylor and Brown, 1988; Myers, 1998). Second, individuals make egocentric comparisons, that is, in order to evaluate the skills of others, they apply the standards that they use to themselves (Dunning et al., 1991; Dunning and Hayes, 1996). Third, the more ambiguous or vague the definition of the skill or trait under evaluation is, the greater the tendency will be to overplace (Felson, 1981; Dunning et al., 1989). Fourth, overplacement increases with the degree to which people claim to be able to control the skill or trait (Alicke, 1985; Dunning, 1993). Fifth, low skill individuals overplace themselves, high skill individuals underplace themselves, and low skill individuals are worse at evaluating their relative skills than high skill individuals—the “unskilled and unaware” hypothesis (Kruger and Dunning, 1999). Sixth, individuals tend to display overplacement on easy tasks and underplacement on hard tasks (Kruger, 1999; Kruger and Dunning, 1999; Krueger and Mueller, 2002; Moore and Kim, 2003; Moore and Healy, 2008). Seventh, most studies find no gender differences in overplacement but some find that men overplace themselves more than women in male-type domains (Ring et al., 2016; Bordalo et al., 2019).

### 2.1.3 Overprecision

Psychology studies on overprecision typically ask participants to provide numerical estimates to general-knowledge questions such as “*What is the height of Mount Everest?*” or numerical forecasts to realizations of random variables such as “*Over the next year, I expect the annual S&P 500 return will be?*”. The participants are then asked to provide estimates of  $x$  percent confidence intervals around their point estimates or forecasts. There is overprecision when the estimated confidence intervals are too narrow, that is, the true answer (8,848 meters) or the ex-post realization of the random variable (next year’s annual S&P 500 return) is inside the participants’  $x$  percent confidence intervals less than  $x$  percent of the time.

Another commonly used task to study overprecision is to ask participants a series of general-knowledge questions in which they have to choose which of two alternatives is correct. For example, “*What is absinthe: a precious stone or a liqueur?*” After participants make their choice, they are asked to state their confidence in having chosen the correct answer, usually on a scale of 50 to 100 percent. There is overprecision when the mean confidence across all questions is greater than the percentage of correct answers.

The main findings of psychology research on overprecision are as follows. First, overprecision is a pervasive finding: typically 90 percent confidence intervals contain the true values only 50 percent of the time and the mean confidence is greater than the percentage of correct answers to general-knowledge questions (Lichtenstein et al., 1982; Russo and Schoe-

maker, 1992; Moore et al., 2015; Hoffrage, 2016). Second, generating reasons why an answer to a general-knowledge question is correct does not lower overprecision, but challenging an answer by asking why it may be wrong lowers overprecision (Koriat et al., 1980). Third, overprecision is lower for easy general-knowledge questions and greater for hard ones (Fischhoff et al., 1977; Lichtenstein et al., 1982; Yates, 1990; Gigerenzer et al., 1991; Griffin and Tversky, 1992). Fourth, underprecision is rarely observed (Moore and Healy, 2008). Fifth, overprecision is correlated with extraversion and narcissism (Schaefer et al., 2004; Campbell et al., 2004). Sixth, when individuals are exposed to repeated feedback about how precise their estimates are, they start by displaying low overprecision but, with minimal feedback, overprecision rapidly spikes to a high level and then levels off (Sanchez and Dunning, 2018).

## 2.2 Evidence from Economics

Economic research on overconfidence differs from psychology research in five main ways. First, it typically provides incentives for truth telling when it elicits beliefs. Second, it often uses observed choices to infer beliefs. Third, it shows that overconfidence affects economic decisions. Fourth, it studies how overconfidence persists after repeated feedback. Fifth, it tests for rational versus irrational overconfidence. DellaVigna (2009) conducts an extensive survey of the empirical evidence on several types of deviations from standard economic theory and overconfidence is discussed as one of three forms of nonstandard beliefs.

### 2.2.1 Eliciting Beliefs with Scoring Rules

Economists typically provide incentives to elicit participants' true beliefs about their skills. This is done with a scoring rule, a function that maps beliefs about a random variable (e.g., one's skill) and the ex-post realization of that random variable (e.g., one's performance on a skill test) into a reward. Schlag et al. (2015) and Schlag and van der Weele (2015) provide excellent reviews of belief elicitation in the laboratory and distinguish two kinds of scoring rules: *deterministic* and *stochastic*.

With a deterministic scoring rule the reward is a fixed prize. The most used and well-known deterministic scoring rule is the quadratic scoring rule (QSR) based on the Brier score (Brier, 1950). With a stochastic scoring rule the reward is a lottery ticket. An example is the binarized scoring rule (BSR) proposed by Hossain and Okui (2013) and Schlag and van der Weele (2013). The main advantage of the BSR over the QSR is that it induces truth-telling independent not only of the participant's utility of wealth, but also of whether the participant is an EU maximizer. This is important given the substantial evidence of heterogeneity in risk preferences (Hey and Orme, 1994; Harless and Camerer, 1994; Starmer, 2000), which suggests that there is a majority of non-EU maximizers and a minority of EU maximizers (Bruhin et al., 2010; Conte et al., 2011; Bruhin et al., 2019). The QSR and the BSR can also be used to measure overprecision (Moore and Healy, 2008; Bruhin et al., 2018). Alternatively, one can ask participants directly for a confidence interval using the Most

Likely Interval (MLI) elicitation rule proposed by Schlag and van der Weele (2015).

There are quite a number of economic experiments on overconfidence that use scoring rules to elicit participants' beliefs. These experiments provide evidence of overconfidence, even in the presence of truth-telling incentives.

Moore and Healy (2008) use the QSR to elicit the probability that participants obtain each possible score, both for themselves and for a randomly selected participant, on easy, hard, and intermediate quizzes. The experiment finds underestimation on easy quizzes, overestimation on hard quizzes, and mean-zero bias on intermediate ones. It also finds overplacement on easy quizzes, underplacement on hard quizzes, and mean-zero bias on intermediate ones. In addition, participants displayed overprecision.

Clark and Friesen (2009) use a stochastic QSR to elicit mean beliefs about absolute and relative performance on two real effort tasks: maximizing the value of a function and decoding five-letter words. The experiment finds mostly unbiased mean estimates of absolute and relative performance.

Park and Santos-Pinto (2010) use the QSR to elicit mean beliefs of poker and chess players' forecasts of relative performance in tournaments. This lab-in-the-field study finds that poker players' forecasts are random guesses with an overplacement bias. Chess players' forecasts are informed guesses, but also display overplacement bias.

Merkle and Weber (2011) use the QSR to elicit participants' entire distribution of beliefs on tests in four different domains: intelligence, memory, general knowledge, and creativity. The experiment finds overplacement in intelligence, memory, and general knowledge (easy domains) but mean-zero bias in creativity (hard domain).

Burks et al. (2013) elicit self-placement of truck drivers on IQ and numeracy tests. Truck drivers name a quintile for their score and are rewarded if their estimate is correct. This elicits the mode of truck drivers' belief distributions and induces truth-telling, independent of the specific form of the utility function, provided the belief distributions are discrete. They find that truck drivers overplace themselves on both tests.

Benoît et al. (2015) use the stochastic scoring rule proposed by Karni (2009) to elicit self-placement in a math and logic test. The experiment finds overplacement: the mean estimate of being placed in the top half is 67 percent, significantly above 50 percent (the mean of an unbiased population).

Bordalo et al. (2019) use a stochastic scoring rule based on the reservation price mechanism of Becker et al. (1964) to elicit the probability that participants answer correctly each of 10 questions, both for themselves and for a randomly selected participant, covering various degrees of difficulty and 12 domains. The experiment finds overestimation of own scores and of others' scores, particularly in more difficult questions and domains. As questions become easier, the extent of overestimation falls, with underestimation for the easiest questions.

## 2.2.2 Inferring Beliefs from Observed Choices

The psychological evidence on overconfidence is mostly based on non-incentivized verbal statements, not on the observation of choice. Some economic experiments use observed choices to make inferences about participants' beliefs, preferences, or welfare. This is one of the basic methods of empirical economics, and is called the revealed-preference approach.

Experiments where participants have to choose between two lotteries (Hoelzl and Rustichini, 2005; Benoît et al., 2015) or between lotteries and certain amounts of money (Murad et al., 2016; Bruhin et al., 2018) provide choice-based measures of beliefs about skill. The main advantage of choice-based measures is that they are more subtle than asking participants directly to state their beliefs. The main disadvantage is that observed choices are often influenced by factors other than beliefs about skill (e.g., risk preferences and strategic sophistication). Hence, inferring beliefs about skill from observed choices requires carefully controlling for additional factors.

Hoelzl and Rustichini (2005) implement a choice-based measure of beliefs about relative skill that relies on a voting game. The experiment finds overplacement on easy skill tests and underplacement on hard skill tests. Furthermore, monetary payments significantly reduce overplacement. The experiment also finds a substantial discrepancy between this choice-based measure of beliefs and stated beliefs, which may be due to the strategic nature of the voting game.

Benoît et al. (2015) study self-placement on a math and logic test using choices between skill-based and luck-based lotteries. The experiment finds overplacement: 52 percent implicitly place themselves in the top 30 percent with a probability of at least .5.

Murad et al. (2016) study overprecision and elicit participants' degree of confidence in their answers, both with and without monetary incentives. The non-incentivized measure generates underprecision on easy questions and overprecision on hard questions. In contrast, the incentivized measure generates underprecision on both types of questions. However, the extent of underprecision is reduced after controlling for utility curvature and probability weighting.

Bruhin et al. (2018) study self-placement and overprecision on an IQ test using two types of beliefs: stated and revealed. Stated beliefs are obtained with a QSR whereas revealed beliefs are inferred from participants' choices via a structural model controlling for utility curvature and probability weighting. The experiment finds overplacement in stated and revealed beliefs and that these are positively correlated, but that the correlation is only .46.

Three main findings can be drawn from these four studies. First, choice-based measures of beliefs also provide evidence for overconfidence. Second, when choice-based measures are used to infer beliefs, the level of overconfidence tends to be lower with incentives for truth telling than without. Third, stated and revealed beliefs are positively correlated but there is a substantial discrepancy between them.

### 2.2.3 Overconfidence and Economic Decisions

A growing literature in experimental economics shows overconfidence influences economic choices such as market entry decisions, innovative activity, entry deterrence, risk taking, and the willingness to compete following success or failure.

In Camerer and Lovo (1999) participants could either enter a market in which the top-ranked entrants would make profits while the bottom-ranked entrants would suffer losses, or stay out. The experiment manipulates recruitment instructions as well as how participants' ranks are determined. In the recruitment condition without self-selection participants are not aware they can earn more if they are good at sports and current events trivia, whereas in the recruitment condition with self-selection this information is provided. In the random condition, participants' ranks are determined randomly whereas in the skill condition ranks are determined according to a trivia quiz. Without self-selection, the average industry profit is positive in the random and skill conditions which implies underentry in both conditions. With self-selection, the average industry profit is positive in the random condition and negative in the skill condition, which implies underentry in the random condition and overentry in the skill condition. Hence, with self-selection, overplacement leads to overentry.

Herz et al. (2014) show that overestimation enhances innovation but overprecision inhibits it. The experiment is based on the management task developed by Ederer and Manso (2013), which captures the trade-off between exploitation and exploration; a core element of innovative decisions. Participants take on the role of a manager at a virtual ice-cream stand, whose aim is to find the product mix that maximizes profit, and are provided with a default business strategy which delivers a known profit. Fine-tuning this strategy (exploitation) allows to increase profits to a certain extent. However, in order to increase profits substantially, participants need to be willing to change the product mix radically (exploration). The experiment finds that overestimation is positively associated with exploration whereas overprecision is negatively linked to exploration.

Cain et al. (2015) investigate entry decisions into one of two markets in which the chance of winning a single prize depends on the number and relative skill of entrants. In the "easy" market the relative skill depends on performance in an easy trivia quiz, while in the "hard" market it depends on performance in a hard trivia quiz. The experiment finds overplacement in the easy quiz, underplacement in the hard quiz, and overentry in the "easy" market.

Bruhin et al. (2018) study how self-placement affects risk taking. The experiment finds that low (high) skill participants are more (less) willing to take risks on lotteries where the winning probabilities depend on relative skill than on lotteries where winning probabilities are exogenous. This pattern of risk taking is driven by systematic biases in beliefs about relative skill: low skill participants overplace themselves, while high skill participants underplace themselves. This finding has important implications for occupational and career choices. Low skill people who overplace themselves might enter professions where performance depends on rela-

tive skill and, as a consequence, crowd out intermediate skill people with correct beliefs.

In Charness et al. (2018) participants are first asked, under a QSR, to state the probability they are in the top 50 percent in an IQ test. Next, participants are randomly matched in pairs. In the baseline treatment, each pair competes in a tournament where the highest score in the IQ test determines the winner. In the strategic treatments, each pair is divided into a sender and a receiver. Receivers can opt out of the tournament after observing the sender’s reported self-confidence, giving senders incentives to distort their reported confidence. There is overplacement in the baseline treatment: the mean estimate of being placed in the top half is 63.4 percent. In the strategic treatments, participants are very likely to enter the tournament when they know that their stated self-confidence is higher than their partner’s, but rarely when the reverse is true. Senders exploit this fact by inflating their reported self-confidence to deter entry.

Huang and Murad (2017) study how relative performance feedback on one task affects self-confidence and the willingness to compete in a second unrelated task. Participants first work on a visual perception task and then add up sets of five randomly generated two-digit numbers. Performance is rewarded by a piece rate in both tasks. In the third task of the experiment, participants’ valuations for submitting their second task performance to a comparative pay scheme are elicited. The experiment shows that those who receive positive feedback in the first task display higher valuations for submitting second task performance to comparative pay than those who receive negative feedback. Similarly, Banerjee et al. (2019) find that success when competing in a first task—memorizing 15 numbers, randomly drawn between 0 and 100, one at a time—increases self-confidence and competitiveness in a subsequent unrelated motor skill task—throwing plastic balls in a bucket.

These studies show overconfidence influences economic decisions such as market entry, innovation, entry deterrence, risk taking, and the willingness to compete. Sections 4.2 and 5.2 review studies which show overconfidence also influences labor market decisions.

#### **2.2.4 Overconfidence and Repeated Feedback**

How do overconfident individuals update their beliefs when faced with information about their skills or performance? Does repeated feedback make overconfidence disappear? Do individuals update beliefs as much as Bayes’ rule predicts? Is updating symmetric with respect to positive and negative feedback? An emerging literature is looking for answers to these questions.

Eil and Rao (2011) find overplacement and asymmetric updating for beauty: participants’ posterior beliefs are more sensitive to positive than to negative feedback. In contrast, the experiment reports unbiased mean estimates and symmetric updating for IQ.

Mobius et al. (2011) find overplacement, conservatism (less updating than predicted by Bayes’ rule), and stronger updating to positive than negative feedback for IQ.

Grossman and Owens (2012) find overestimation of absolute performance in four different multiple-choice quizzes. While feedback improves absolute performance estimates, this learning does not translate into improved estimates of subsequent performance. There is no evidence of asymmetric updating about absolute performance.

Wiswall and Zafar (2015) find that college students at New York University are substantially misinformed about population earnings by academic major. There is also evidence of asymmetric updating with stronger updating about personal future earnings after positive information about population earnings than after negative information.

Zimmermann (2019) elicits beliefs about rank on an IQ test after noisy feedback. In one treatment, beliefs are elicited immediately after the feedback, while in a second treatment beliefs are elicited one month after. The study finds positive feedback has a persistent effect on beliefs. In contrast, negative feedback only influences beliefs in the short-run. This is due to an asymmetry in the recall of feedback (a form of selective memory).

Huffman et al. (2019) find that managers of a chain of food-and-beverage stores who compete repeatedly in high-stakes tournaments overplace themselves relative to a range of different predictors obtained from past tournament outcomes. Overplacement is persistent under repeated feedback and there is evidence of selective memory: managers with poorer past performances have larger recall errors, and these are skewed towards overly positive memories. In addition, managers who have overly-positive memories of past feedback are those who are particularly likely to overplace themselves.

Three main conclusions can be drawn from this literature. First, overconfidence can persist after repeated feedback even with incentives for truth telling. Second, many studies find evidence for conservatism and selective memory. Third, evidence for asymmetric Bayesian updating is mixed: some studies find stronger updating after positive than negative feedback (Eil and Rao, 2011; Mobius et al., 2011; Wiswall and Zafar, 2015), some find the opposite (Coutts, 2019), and others find symmetric updating (Eil and Rao, 2011; Grossman and Owens, 2012; Buser et al., 2018).

## 2.2.5 Rational versus Irrational Overconfidence

The psychology research does not distinguish rational from irrational overconfidence. If individuals have imperfect knowledge of their ability, then rational overconfidence can be generated by Bayesian updating from a common prior (and truthful revelation of posteriors) whereas irrational overconfidence cannot.

Benoît and Dubra (2011) show that Bayesian updating from a common prior imposes testable restrictions on the distribution of relative ability judgments. These implications allow researchers to distinguish rational from irrational overplacement. For example, when drivers are asked to state the probability of being placed in the top half of the distribution of skill, the mean estimate of being placed in the top half should be 50 percent; the mean of an unbiased population. Additionally, when individuals

are asked to place themselves into deciles, up to 2/10 of the people can rationally place themselves in the top decile, up to 4/10 can rationally place themselves in the two top deciles, and up to  $2i/10$  can rationally rank themselves in the top  $i$  deciles, for  $i = 3, 4, 5$ . Bayesian updating also imposes testable restrictions on the joint distribution of relative ability judgments and true abilities. In particular, Burks et al. (2013) show that it must be true that of all individuals placing themselves in ability quantile  $k$ , the largest (modal) share of them must actually be from quantile  $k$ .

The economic evidence on rational versus irrational overplacement is as follows. Hoelzl and Rustichini (2005) and Moore and Healy (2008) find rational overplacement on easy quizzes and rational underplacement on hard quizzes. Niederle and Vesterlund (2007) find irrational overplacement by men and rational overplacement by women on an easy real effort task. Clark and Friesen (2009) find either unbiased mean estimates or underplacement in two hard real effort tasks. Park and Santos-Pinto (2010) find rational overplacement of poker and chess players. Merkle and Weber (2011) find irrational overplacement in easy tests of intelligence, memory, and general knowledge but unbiased mean estimates in a hard test of creativity. Eil and Rao (2011) find rational overplacement in beauty and unbiased mean estimates on an IQ test. Mobius et al. (2011), Burks et al. (2013), Benoît et al. (2015), Charness et al. (2018), and Huffman et al. (2019) find irrational overplacement.

Overall, three main results emerge from these studies. First, there are many instances of irrational overplacement even with incentives for truth telling. Second, easy tasks often lead to irrational overplacement, intermediate tasks to rational overplacement, and hard tasks to rational underplacement. Third, none of the studies reports irrational underplacement.

### 3 Overconfidence: Explanations

This section reviews economic explanations for overconfidence. These are divided into five main types: Bayesian updating from a common prior, differing priors or opinions, biases in information processing, strategic benefits, and non-strategic benefits. Psychologists have two main types of explanations for overconfidence: cognitive and motivational (Alicke and Govorun, 2005). According to cognitive explanations, overconfidence is due to information-processing mechanisms that are benefit-neutral (Nisbett and Ross, 1980). According to motivational explanations, overconfidence is driven by a conscious or unconscious desire to arrive at a particular positive belief about one's ability (Kunda, 1990). Economic explanations such as Bayesian updating from a common prior, differing priors or opinions, and biases in information processing fit the category of cognitive explanations. In contrast, economic explanations based on strategic or non-strategic benefits fit the category of motivational explanations.

### 3.1 Bayesian Updating from a Common Prior

Overconfidence can be the result of Bayesian updating from a common prior (Zábojník, 2004; Benoît and Dubra, 2011). This type of overconfidence is rational and does not lead to ex-ante costly decision-making mistakes.

In Zábojník (2004) agents are uncertain about their ability and can undertake experiments to learn it. Agents start out with a common prior belief and update beliefs about their ability from the signals they observe using Bayes' rule. The costs of experimenting are proportional to expected output which increases in expected ability. Agents will continue testing their abilities until their posterior beliefs become high enough, at which point they stop. Those with higher beliefs start producing early, since their opportunity cost of experimenting is higher. In contrast, agents with lower beliefs keep experimenting until they strike a string of good signals, and so will end up with high posteriors. This way, the share of agents with high posterior beliefs grows over time.

In Benoît and Dubra (2011) agents are uncertain about their ability (or type) and receive information about it from their personal experience. The population starts out with a common prior belief and agents update beliefs about their ability from the signals they observe using Bayes' rule. If unfavorable signals are rare (the task is easy), there is rational overplacement. In contrast, if unfavorable signals are frequent (the task is hard), there is rational underplacement. Over time, as signals accumulate, the agents' posterior beliefs converge to their true types and the population ends up with correct beliefs.

### 3.2 Differing Priors or Opinions

Irrational overconfidence is incompatible with Bayesian updating from a common prior. However, Bayesian rational agents do not need to hold common prior beliefs (Morris, 1995; Gul, 1998). Van den Steen (2004) and Santos-Pinto and Sobel (2005) show irrational overplacement can arise naturally in a population of Bayesian rational agents with differing priors or opinions. Van den Steen (2011) shows that relaxing the common prior assumption can also explain overprecision.

In Van den Steen (2004) agents must choose from a finite number of actions (e.g., different driving styles) and have different prior beliefs about the probability that a given action will be successful (e.g., avoiding an accident). Since each agent selects the action that (in his view) is most likely to succeed, each agent believes that his choice is at least as good as the choices made by others in the population, and everyone believes that the other individuals overestimate their probability of succeeding. The population displays irrational overplacement since each agent expects to do better than all other agents.

In Santos-Pinto and Sobel (2005) agents begin with an initial endowment of skills and make skill investments to maximize their ability subject to a budget constraint. There is a subjective component to the definition of ability, in that different agents hold different opinions about how skills (e.g., parking, overtaking, driving in heavy city traffic, driving in a snow

storm, etc.) combine to determine an ability level (e.g., driving). When an agent responds to questions about relative standing, he responds egocentrically, that is, he uses his opinion to compare his final skills to those of others in the population. Irrational overplacement arises because agents tailor their skill augmentation to their own opinion and use it to evaluate others' final skills. Much of the evidence about overplacement fits comfortably in this model.

In Van den Steen (2011) agents start with a common prior about the mean of a random variable, observe the same identical signals, but have differing priors about the variance of signals. In trying to find the best estimate, agents will place more weight on signals they subjectively consider to be more precise. As a consequence, they put too much weight on signals for which they overestimate the variance and not enough weight on signals for which they underestimate the variance. This generates overprecision.

### 3.3 Biases in Information Processing

Overconfidence can be a consequence of biases in information processing such as the confirmation bias (Rabin and Schrag, 1999) or the self-serving bias in causal attributions (Gervais and Odean, 2001). These two explanations maintain the common prior assumption but relax Bayesian updating. This type of overconfidence leads to costly errors in decision-making.

In Rabin and Schrag (1999) there are two possible states of the world and agents receive binary signals that are correlated with the true state. Agents initially view the two states as equally likely and, after receiving each signal, update their beliefs about the true state. Rabin and Schrag (1999) model the confirmation bias (Lord et al., 1979; Klayman and Ha, 1987) by assuming that when an agent receives a signal that is counter to his current belief about which state is more likely, there is a positive probability that he misinterprets that signal. The model shows that the first signals an agent observes play a disproportional large role in determining his posterior beliefs and that the agent displays overprecision in the sense that his belief in favor of one state is stronger than what is justified by the available evidence.

In Gervais and Odean (2001) agents start with a common prior belief about their ability, observe a sequence of signals, and display a learning bias inspired by the self-serving bias in causal attributions: the fact that people tend to attribute success to skill and failure to bad luck (Miller and Ross, 1975). Gervais and Odean (2001) model this learning bias by assuming agents overweight their successes when they form their posterior beliefs. The model shows that as soon as an agent observes one success, he overestimates his ability. In the short run, after a few signals, agents will tend to overestimate their abilities. In the long run, as signals accumulate and provided that the learning bias is not too large, the population ends up with correct beliefs.

## 3.4 Strategic Benefits

Overconfidence can exist because it provides strategic benefits that compensate for its decision-making costs. These strategic benefits can be either interpersonal or intrapersonal. In the former case, overconfidence changes a rival's behavior to the benefit of an overconfident agent (Trivers, 1985; Heifetz et al., 2007a,b; Von Hippel and Trivers, 2011). In the latter case, overconfidence changes the behavior of one's future self to the benefit of one's present self (Bénabou and Tirole, 2002).

### 3.4.1 Overconfidence as Commitment Device

In Heifetz et al. (2007a) a large population of agents are continuously and randomly matched in pairs to interact with one another. Agents may differ in the way they perceive the returns of their actions. An overconfident agent overestimates the return to his action for any given action taken by the rival while an underconfident agent underestimates it. Agents' perceptions are perfectly observable. In every pairwise interaction, the matched agents choose actions to maximize their perceived payoff functions and receive payoffs according to their actual payoff functions. Actions can be either strategic substitutes or complements. The proportion of more successful perceptions in the population increases over time at the expense of less successful perceptions. Heifetz et al. (2007a) show that the distribution of perceptions converges to a unit mass where agents slightly overestimate the returns to their actions. All other perceptions, including correct ones, become extinct asymptotically. The intuition is as follows. Overconfident agents play more aggressively than realists or underconfident agents, as they exaggerate the impact of their actions on their payoffs. Being aggressive has a direct cost because an overconfident agent fails to play a best-response against the rival's action. However, being aggressive also has a strategic benefit. When actions are strategic substitutes (complements), the aggressive behavior of an overconfident agent induces the rival to play soft (aggressively) which benefits the overconfident agent. When the agent is slightly overconfident, the negative direct effect is of second-order, as it results from a small deviation from payoff-maximizing behavior, whereas the strategic effect resulting from the change in the rival's action is of first-order. Hence, slightly overconfident agents will do better than significantly overconfident agents and will gradually take over the whole population. Heifetz et al. (2007b) extend this result to almost every strategic interaction and even to those in which agents' perceptions are only imperfectly observed. Overall, the assumption that decision-makers must hold correct beliefs about their skills cannot be justified by appealing to evolutionary arguments.

### 3.4.2 Self-Deception to Better Deceive Others

Overestimation of ability might lower the cognitive or moral cost of deceiving others when it is easier to convince others that one has high ability, if one believes this as well (Trivers, 1985; Von Hippel and Trivers, 2011). Even though this explanation is from evolutionary biology and not from economics, it is included here since experimental economic studies find

support for it. On the one hand, Gneezy (2005), Gneezy et al. (2018), and many others, find that individuals suffer cognitive or moral costs when they lie to others. On the other hand, Schwardmann and Van der Weele (2019) and Solda et al. (2019) provide causal evidence that people deceive themselves in order to more effectively deceive others and demonstrate that this is an effective persuasion strategy.

### 3.4.3 Strategic Information Avoidance

Bénabou and Tirole (2002) show that a sophisticated agent with time-inconsistent and present-biased preferences might strategically prefer to avoid information about ability to remain overconfident. The agent is risk neutral, is uncertain about his ability, and displays quasi-hyperbolic discounting (Strotz, 1955; Phelps and Pollak, 1968). The agent cannot commit to his future decisions but is sophisticated, that is, he can fully anticipate his future behavior and take actions to influence it (Laibson, 1997; O'Donoghue and Rabin, 1999; Carrillo and Mariotti, 2000). In the first period, self 0 starts out with a prior belief about ability and can either learn ability at zero cost or remain uninformed. In the second period, self 1 can either undertake an activity or not. Effort and ability are complements in production and undertaking the activity leads to a cost in the second period and a delayed reward in the third period. Bénabou and Tirole (2002) establish that if the agent starts out with a high enough prior belief (or self-confidence), then self 0 prefers to remain uninformed about ability. The intuition is as follows. Avoiding information to maintain a high self-confidence has a benefit and a cost for self 0. Since effort and ability are complements, high self-confidence raises self 1's motivation to undertake the activity. This is beneficial for self 0 as it counteracts self 1's tendency to slack for intermediate values of ability. However, if the agent's ability is low but he is unaware of it, then high self-confidence leads self 1 to undertake the activity. This is harmful for self 0 as self 0 would prefer that self 1 does not undertake the activity when ability is low. When the agent's self-confidence is high enough, the benefit is larger than the cost, and self 0 strategically prefers to avoid information to maintain a high self-confidence thereby offsetting self 1's lack of motivation to undertake the activity. Eil and Rao (2011) find experimental support for this explanation: participants who receive negative feedback are less likely to choose to costlessly obtain full information about their rank, and some are even willing to pay to avoid obtaining it.

### 3.4.4 Selective Memory

Bénabou and Tirole (2002) propose an alternative mechanism for overestimation of ability: selective memory. In the first period, self 0 receives a binary signal about his ability (bad or no news) and then decides the recall probability of bad news. In the second period, self 1 chooses between undertaking the activity or not assessing the credibility of memory about ability using Bayes' rule and self 0's choice of recall probability. Faced with a signal that hurts his self-confidence, self 0 chooses the recall probability so as to maximize the difference between his expected benefit

and his memory cost. Bénabou and Tirole (2002) establish that for high degrees of time-inconsistency, there exists a unique equilibrium with maximum repression of bad news. Chew et al. (2018), building on Bénabou and Tirole (2002), consider signals with three realizations (good, bad, and no news) and assume the agent derives psychological benefits from holding a positive self-image. This model is able to explain false memory in favor of good news as well as selective amnesia against bad news. Chew et al. (2018), Zimmermann (2019), and Huffman et al. (2019) provide experimental support for selective memory.

### 3.5 Non-Strategic Benefits

Overconfidence might exist because it provides non-strategic benefits that compensate for its decision-making costs. For instance, the genetic advantage from increased offspring (Waldman, 1994), the material benefits of performing better at a task (Compte and Postlewaite, 2004), or the psychological benefits from ego utility (Kőszegi, 2006), or from social approval (Bénabou and Tirole, 2006; Ewers and Zimmermann, 2015).

Waldman (1994) shows overestimation of ability by males can be an evolutionarily stable outcome given sexual inheritance and complementarity between effort and self-confidence. Compte and Postlewaite (2004) show that overestimation of the probability of success at a task can be helpful if an agent's future chances of success at that task depend on self-confidence (e.g., tasks where nervousness or stress are detrimental to performance).

Kőszegi (2006) shows that overestimation of ability can arise when individuals have ego utility or the desire to hold a positive self-image regardless of others' judgments or reactions. When this is the case, an individual will use his actions to manipulate his self-image. In Kőszegi (2006) agents feel better with a favorable self-perception, even at the cost of being overconfident and thus making wrong choices.

Bénabou and Tirole (2006) and Ewers and Zimmermann (2015) show that overestimation of ability can also arise when an individual derives utility from others having favorable views of his knowledge or skills (a desire for social approval). When this is the case, an individual will use his actions in a social setting to manipulate others' judgements. Ewers and Zimmermann (2015) test this explanation experimentally and find that the desire to impress an audience causes participants to state higher self-confidence. However, as the chapter shows, much of the evidence on overconfidence has been obtained in settings without an audience which rules out social approval motives.

## 4 Internal Labor Markets

This section reviews research on the impact of worker overconfidence on internal labor markets.

## 4.1 Theory

Section 4.1 surveys theory models on how employee overconfidence affects labor markets where wages are determined by bargaining between workers and firms. Kőszegi (2014) offers a literature review of contract theory when agents have mistaken beliefs in general, and discusses some of the implications of overconfidence for contracting.

### 4.1.1 Basic Model

This section presents a principal-agent model that illustrates the main implications of worker overconfidence for contracting. The most basic model that captures the problem of motivating a worker to exert effort has two actions and two outcomes (Laffont and Martimort, 2002). Consider a principal (she) who hires an overconfident agent (he) to perform a project. The outcome of the project can be either high,  $q_H$ , or low,  $q_L$ , with  $q_H > q_L > 0$ . The outcome of the project is observable and verifiable. The agent can either exert high effort,  $e = 1$ , or low effort,  $e = 0$ . The agent is overconfident in the sense that he overestimates the returns to effort. If the agent exerts high effort, the true probability of a high outcome is  $p_1$ , but the agent has an incorrect belief in that he believes the probability of a high outcome is  $\tilde{p}_1 \in (p_1, 1)$ . If the agent exerts low effort, the true probability of a high outcome is  $p_0 \in [0, p_1)$ , but the agent has an incorrect belief in that he believes the probability of a high outcome is  $\tilde{p}_0 \in (p_0, \tilde{p}_1)$ . Note that an overconfident agent still perceives high effort to be more productive than low effort since  $\tilde{p}_1 > \tilde{p}_0$ .

The principal is risk neutral and cares only about maximizing expected profits: the difference between expected benefits and expected wage. The principal is perfectly informed about the agent's true productivity and his mistaken beliefs. The principal and agent know each other's beliefs and they both believe the other to be mistaken, that is, they "agree to disagree." The principal makes a take-it-or-leave-it contract offer to the agent which promises to pay wage  $w_H$  when the outcome of the project is  $q_H$  and  $w_L$  when it is  $q_L$  with  $w_H \geq w_L$ . The agent can either accept or reject the contract. If the agent accepts the contract and chooses effort  $e$ , the principal's expected profits are then  $E(\Pi) = E(Q) - E(W)$ , where  $E(Q) = p_e q_H + (1 - p_e) q_L$  is the expected outcome of the project and  $E(W) = p_e w_H + (1 - p_e) w_L$  is the expected wage. The gap between  $q_H$  and  $q_L$  is assumed to be high enough such that the principal always prefers to implement high effort.

The agent is risk averse with a utility function that is separable in wage and effort  $U(w, e) = u(w) - c(e)$ , with  $u$  strictly increasing and concave and  $c$  increasing. The agent's cost of effort is  $c(e) = ce$  with  $c > 0$ . Hence, the cost of high effort is normalized to  $c(1) = c$  and the cost of low effort to zero  $c(0) = 0$ . The agent's perceived expected utility from accepting the contract and choosing effort  $e$  is  $E_{\tilde{p}_e}(U) = \tilde{p}_e u(w_H) + (1 - \tilde{p}_e) u(w_L) - ce$ . This must be greater than or equal to his reservation utility level  $\underline{u}$ . The agent holds correct beliefs about his cost of effort and his reservation utility. Taking the perspective of an outside observer who knows the true productivity of the agent, the actual expected utility of an overconfident

agent who chooses effort  $e$  is  $E_{p_e}(U) = p_e u(w_H) + (1 - p_e)u(w_L) - ce$ , which differs from his perceived expected utility.

If effort is observable and verifiable in a court of law, the principal's contract offer specifies that the agent must exert high effort. Therefore, the principal's problem is to implement high effort at the lowest possible expected wage:

$$\begin{aligned} \min_{w_H, w_L} \quad & p_1 w_H + (1 - p_1)w_L \\ \text{s.t.} \quad & \tilde{p}_1 u(w_H) + (1 - \tilde{p}_1)u(w_L) - c \geq \underline{u}. \end{aligned} \quad (1)$$

The constraint in (1) is the agent's participation constraint or PC. For the agent to accept the contract his perceived expected utility of exerting high effort must be no less than his reservation utility  $\underline{u}$ . The solution to (1) is characterized by the conditions

$$\frac{p_1}{1 - p_1} = \frac{\tilde{p}_1}{1 - \tilde{p}_1} \frac{u'(w_H)}{u'(w_L)}, \quad (2)$$

and

$$\tilde{p}_1 u(w_H) + (1 - \tilde{p}_1)u(w_L) - c = \underline{u}. \quad (3)$$

In the absence of overconfidence, such that  $\tilde{p}_1 = p_1$ , the risk-neutral principal would fully insure a risk-averse agent, a well-known application of the Borch rule (Borch, 1962). Since  $\tilde{p}_1 > p_1$ , it follows from (2) that  $u'(w_H) < u'(w_L)$  which implies  $w_H > w_L$ . Hence, when effort is observable, a risk-neutral principal does not fully insure an overconfident and risk-averse agent. The intuition is as follows. When the agent overestimates the probability the high outcome will be realized, the principal can decrease her expected wage by promising to pay a bonus when the outcome of the project is high and a penalty when the outcome of the project is low. The principal and agent are *de facto* wagering on the outcome of the project. The agent mistakenly views this wager as generating a positive expected value which makes him willing to accept some level of risk. For a small enough deviation from full insurance, the agent's overconfidence generates a first-order gain to the principal, whereas risk aversion generates a second-order loss. From the perspective of an outside observer, the agent is worse off accepting the principal's wager since  $\tilde{p}_1 > p_1$  implies that his actual expected utility is lower than his perceived expected utility. In sum, when effort is observable, the principal exploits the agent's overconfidence by offering him a contract with high-powered incentives which exposes the agent to risk.

If effort is unobservable or unverifiable, the contract offer must be such that the agent accepts it and finds it in his best interest to exert high effort. The principal's problem thus becomes

$$\begin{aligned} \min_{w_H, w_L} \quad & p_1 w_H + (1 - p_1)w_L \\ \text{s.t.} \quad & \tilde{p}_1 u(w_H) + (1 - \tilde{p}_1)u(w_L) - c \geq \underline{u} \\ & \tilde{p}_1 u(w_H) + (1 - \tilde{p}_1)u(w_L) - c \geq \tilde{p}_0 u(w_H) + (1 - \tilde{p}_0)u(w_L). \end{aligned} \quad (4)$$

The first constraint in (4) is the agent's PC and the second constraint is the agent's incentive constraint or IC. In any solution to (4) the agent's participation constraint must hold with equality, otherwise the principal could reduce her agency cost by lowering  $w_H$  and  $w_L$  while keeping the power of incentives, that is, the utility gap  $u(w_H) - u(w_L)$ , constant. Hence, as in the case when effort was observable, an overconfident agent is always worse off than an agent with correct beliefs.

Santos-Pinto (2008) distinguishes between two effects the agent's overconfidence has on the principal's problem when effort is unobservable. The *participation effect* is the impact of overconfidence on the agent's willingness to accept the contract offered by the principal for a fixed compensation scheme. The *incentive effect* is the impact of overconfidence on the agent's willingness to exert high effort for a fixed compensation scheme.

The participation effect of overconfidence is always favorable to the principal. To see this, start by recalling the features of the optimal solution to the principal's problem when effort is unobservable and the agent has correct beliefs, that is,  $\tilde{p}_1 = p_1$  and  $\tilde{p}_0 = p_0$ . In this case both constraints are satisfied with equality and there is a trade-off between incentives and insurance. High effort can only be implemented if the principal exposes the agent to risk by offering him a high wage when the outcome of the project is high and a low wage when it is low (one can interpret this as a base wage plus a bonus). An overconfident agent thinks he is more likely to attain a high outcome, and hence to earn a high wage, than an agent with correct beliefs. Hence, (increased) agent overconfidence has a favorable participation effect (relaxes the PC).

The incentive effect of overconfidence can be either favorable or unfavorable to the principal. To see this rewrite IC as

$$u(w_H) - u(w_L) \geq \frac{c}{\Delta\tilde{p}}, \quad (5)$$

where  $\Delta\tilde{p} \equiv \tilde{p}_1 - \tilde{p}_0$ . It follows from (5) that agent overconfidence has a favorable incentive effect (relaxes the IC), when  $\Delta\tilde{p} > \Delta p \equiv p_1 - p_0$ . Note that  $\Delta\tilde{p} > \Delta p$  is equivalent to  $\tilde{p}_1 - p_1 > \tilde{p}_0 - p_0$ , that is, the agent overestimates the productivity of high effort more than he overestimates the productivity of low effort. In this case effort and self-confidence are *complements*: for a fixed compensation scheme, an overconfident agent is more willing to exert high effort than an agent with correct beliefs. Complementarity of effort and self-confidence occurs when effort and ability are complements in the production of performance, a usual assumption. In contrast, overconfidence has an unfavorable incentive effect when  $\Delta\tilde{p} < \Delta p$ . In other words, overconfidence lowers the agent's willingness to exert high effort when the agent overestimates the productivity of low effort more than he overestimates the productivity of high effort. In this case, effort and self-confidence are *substitutes*: for a fixed compensation scheme, an overconfident agent is less willing to exert high effort than an agent with correct beliefs. Substitutability of effort and self-confidence occurs when effort and ability are substitutes in the production of performance which may occur, for example, when the agent works towards a fixed quota or goal.

Overconfidence can thus lead to two different situations in terms of

the parties' welfare when effort is unobservable. First, overconfidence can make the principal better off and the agent worse off. In this situation, the principal exploits the agent's overconfidence since she is able to pay a lower expected wage to an overconfident agent than to an agent with correct beliefs. This happens when either (i) effort and self-confidence are complements or (ii) effort and self-confidence are substitutes and the unfavorable incentive effect is small. Second, overconfidence can make both parties worse off. In this situation, the principal pays a higher expected wage to an overconfident agent than to an agent with correct beliefs. This happens when effort and self-confidence are substitutes and the unfavorable incentive effect is large.

de la Rosa (2011) shows the agent's degree of overconfidence determines not only which of the two situations described above arise, but also the relationship between overconfidence and the contract's power of incentives. More precisely, it depends on whether the agent is *slightly* or *significantly* overconfident. Formally, the agent is said to be slightly overconfident if

$$\frac{p_1}{1-p_1} \geq \frac{\tilde{p}_1}{1-\tilde{p}_1} \frac{u'(\bar{w}_H)}{u'(\bar{w}_L)},$$

conversely, the agent is said to be significantly overconfident if

$$\frac{p_1}{1-p_1} < \frac{\tilde{p}_1}{1-\tilde{p}_1} \frac{u'(\bar{w}_H)}{u'(\bar{w}_L)},$$

where  $\{\bar{w}_H, \bar{w}_L\}$  is the contract that satisfies the agent's PC and IC with equality, that is,  $\{\bar{w}_H, \bar{w}_L\}$  satisfies  $u(\bar{w}_H) = \underline{u} + (1-\tilde{p}_0)c/\Delta\tilde{p}$  and  $u(\bar{w}_L) = \underline{u} - \tilde{p}_0c/\Delta\tilde{p}$ .

When the agent is slightly overconfident, IC binds and the solution to (4) is given by  $\{\bar{w}_H, \bar{w}_L\}$ . In this case, the principal's main consideration is to provide incentives for the agent to exert high effort, not wagering against him, and incentive provision defines the contract's power of incentives, which is  $u(\bar{w}_H) - u(\bar{w}_L) = c/\Delta\tilde{p}$ . If effort and self-confidence are complements,  $\Delta\tilde{p} > \Delta p$ , the principal *lowers* the power of incentives and the agent bears less risk, which reduces the principal's agency cost. Hence, when effort is unobservable, the agent is slightly overconfident, and effort and self-confidence are complements, the principal is better off with an overconfident agent. However, if effort and self-confidence are substitutes,  $\Delta\tilde{p} < \Delta p$ , the principal *raises* the power of incentives and the agent is exposed to more risk, which increases the principal's agency cost. Hence, when effort is unobservable, the agent is slightly overconfident, and effort and self-confidence are substitutes, the principal is better (worse) off with an overconfident agent when the unfavorable incentive effect is small (large).

When the agent is significantly overconfident, IC is not binding and the solution to (4) is given by (2) and (3). In this case, the principal's main consideration is exploiting the agent's mistaken beliefs and the wagering as a consequence of disagreement, implied in (2), defines the contract's power of incentives. Hence, when effort is unobservable and the agent is significantly overconfident, the principal is better off with an overconfident agent than with one with correct beliefs.

While most individuals are overconfident, some are underconfident. The basic model can also be used to illustrate the implications of underconfidence for contracting. Assume the principal still wishes to implement high effort but now she faces an underconfident agent, that is, an agent with beliefs  $\tilde{p}_0 \in (0, p_0)$  and  $\tilde{p}_1 \in (\tilde{p}_0, p_1)$ . When effort is observable, the principal wagers against an underconfident agent by offering him a contract with negative incentives, that is, a contract that offers a high wage when the outcome of the project is low and a low wage when it is high (since  $\tilde{p}_1 < p_1$ , it follows from (2) that  $u'(w_H) > u'(w_L)$  which, in turn, implies  $w_H < w_L$ ). The principal is better off with this wager and, from the perspective of an outside observer, the underconfident agent is worse off.

When effort is unobservable and the agent is underconfident, the participation effect of underconfidence is unfavorable to the principal. High effort can only be implemented if the principal exposes the agent to risk by offering him a high wage when the outcome of the project is high and a low wage when it is low. However, an underconfident agent thinks he is more likely to attain a low outcome, and hence to earn a low wage, than an agent with correct beliefs. Thus, (increased) agent underconfidence has an unfavorable participation effect (worsens the PC). The incentive effect of underconfidence, like that of overconfidence, can be either favorable or unfavorable to the principal. From this it follows that the principal tends to be worse off with an underconfident agent. The only exception is when the incentive effect of underconfidence is favorable and large.

In sum, the basic model shows overconfident workers tend to be more desirable employees than workers with correct beliefs. Generally, the firm can exploit an overconfident worker by appropriately adjusting her compensation contract and paying him less than what a worker with correct beliefs would demand. Overconfident workers also tend to be more desirable employees than underconfident workers. The reason is that, generally, a firm cannot exploit an underconfident worker when effort is unobservable. Finally, the basic model shows worker overconfidence may result in compensation contracts with lower or higher power of incentives, depending both on the level of overconfidence and the relation between effort and self-confidence. One important implication of these results is that firms will treat overconfident workers differently from other workers. This might lead to systematic differences in labor market outcomes among groups which vary in self-confidence or which vary in attributes informative about self-confidence.

#### 4.1.2 Competing Principals

The basic model represents labor market situations where the employer has all the bargaining power in the relationship. However, in labor market situations where many employers compete to hire a few qualified employees, the bargaining power is in the hands of the employees (e.g., sports, music, and arts superstars). Contrary to the basic model, overconfidence can now lead to a Pareto improvement.

Following de la Rosa (2008), suppose multiple principals compete to hire a single risk averse agent and share the same correct beliefs about the

agent's productivity. The timing of the interaction is as follows. First, the principals make simultaneous contract offers to the agent. Second, the agent chooses which contract (if any) to accept. Third, if the agent accepts the contract he then chooses effort. Fourth, the outcome of the project is realized and payoffs are determined according to the contract. The outside option for each principal is not contracting with the agent which yields zero profits.

An equilibrium contract is such that no other contract can (i) attract the agent by offering him terms that he strictly prefers and (ii) yield higher expected profits for the offering principal. In other words, the principals compete for the agent by maximizing the agent's perceived expected utility under a non-negative expected profits condition and the agent's IC. The agent's reservation utility level  $\underline{u}$  is assumed to be low enough such that the agent accepts the contract. Hence, if effort is unobservable, the problem of the principal whose contract offer is accepted by the agent is

$$\max_{w_H, w_L} \tilde{p}_1 u(w_H) + (1 - \tilde{p}_1) u(w_L) - c \quad (6)$$

$$\begin{aligned} \text{s.t. } p_1 w_H + (1 - p_1) w_L &\leq p_1 q_H + (1 - p_1) q_L \\ \tilde{p}_0 u(w_H) + (1 - \tilde{p}_0) u(w_L) &\leq \tilde{p}_1 u(w_H) + (1 - \tilde{p}_1) u(w_L) - c. \end{aligned}$$

The first constraint in (6) is the principal's PC and the second constraint is the agent's IC. In any solution to (6), the principal's PC must hold with equality. The intuition is as follows. Because the principals share the same correct beliefs about the agent's productivity, the expected profits for the principal whose offer is accepted by the agent in equilibrium must be zero. Intuitively, if the offering principal made positive expected profits, another principal could outbid that contract offer by promising to give the agent a slightly higher perceived expected utility while implementing the same effort, thus attracting the agent and making positive expected profits. Since the offering principal makes zero expected profits, social welfare depends only on the agent's expected utility. This implies that the expected wage of an overconfident agent is equal to that of an agent with correct beliefs. Hence, when overconfidence lowers the power of incentives, the overconfident agent is exposed to less risk than an agent with correct beliefs and there is a Pareto improvement.

### 4.1.3 Interdependent versus Individualistic Compensation Schemes

The basic model focuses on how a firm contracts with a single worker. However, most firms employ multiple workers. Santos-Pinto (2008) shows that overplacement makes interdependent compensation schemes (e.g., tournaments) more attractive to firms than individualistic ones (e.g., piece rates). If the principal knows the agents underestimate their peers' productivity, then the principal can make an advantageous wager. She does this by increasing compensation in outcome pairs where an agent's peer has a low outcome and reducing compensation in outcome pairs where an agent's peer has a high output by comparison with the optimal individualistic compensation scheme.

#### 4.1.4 Teamwork

Teamwork is one of the most widespread organizational forms with multiple workers. Under teamwork, the marginal productivity of an agent's effort increases with the efforts of the other agents, that is, agents' effort choices are strategic complements. The classic model of teamwork developed by Holmstrom (1982) suggests that free-rider problems are pervasive in teams. Typically, under teamwork, the agents' individual contributions are neither observable to the firm nor to each other. As a result, agents who shirk cannot be identified. This free-rider problem will generally lead to underprovision of effort in teams.

Gervais and Goldstein (2007) show that overconfidence can not only mitigate free-riding in teams, but also lead to a Pareto improvement. The model consists of an all equity firm, owned by risk neutral shareholders (the principal), and requiring the effort of two agents for production. The team outcome derives from a single project which can either succeed or fail with positive probability. The project generates a positive outcome if it succeeds, and zero if it fails. The probability of success of the project increases with the effort of each agent, with the skill of each agent, and with an interaction term which captures the synergies between the effort choices of the two agents.

Agents are risk neutral and protected by limited liability, that is, all contractual transfers from the firm to the agents must be non-negative. Their effort decisions are made simultaneously and are unobservable to the other agent and to the firm. Since effort is non-contractible, contracts offered to the agents must specify how much each agent receives when the project succeeds and when it fails. Agents choose effort to maximize their perceived expected utility and they sustain a utility cost of effort. The agents' perceived expected utilities from accepting the contract must be at least their reservation utility level. Agent 1 has correct beliefs about his skill whereas agent 2 is overconfident. The principal and agent 1 know that agent 2 is overconfident, and agent 2 knows this is what they believe but disagrees with them.

Gervais and Goldstein (2007) start by analyzing the agents' equilibrium effort choices for fixed compensation schemes. Since effort and self-confidence are complements, the higher the overconfidence of agent 2 is the harder agent 2 works because he overestimates the probability the project will succeed. As agent 2's overconfidence increases, agent 1 also works harder because, knowing that agent 2 works harder, he knows that the potential synergies from their combined efforts are higher. Hence, overconfidence can mitigate free-riding in teams. Gervais and Goldstein (2007) also show that agent 2's overconfidence leads to a Pareto improvement as long as it is not too high; in that case, overconfidence raises firm value, agent 1's expected utility, and agent 2's actual expected utility. Finally, Gervais and Goldstein (2007) show that when the project succeeds, the overconfident agent is paid more than the agent with correct beliefs. Under teamwork, the firm will pay more to an overconfident agent than to an agent with correct beliefs, since an increase in the compensation of the overconfident agent has a greater impact on raising the agents' effort provision than the same increase in the compensation of the agent with

correct beliefs.

#### 4.1.5 Tournaments

Tournaments are one of the most common forms of compensation scheme used in firms with multiple workers. For example, managers are involved in promotion tournaments to become CEOs and salespeople are often paid bonuses that depend on their sales relative to those of other salespeople in the firm. Formally, a tournament is an interdependent compensation scheme with a defined prize structure. The agents make their choices simultaneously and independently, taking the prize structure as given. The agents' actions (e.g., effort, risk exposure) are neither observable to the firm nor to each other. The firm observes the agents' rank in terms of output and awards the prizes to agents according to their rank, with high-rank agents receiving higher prizes than low-rank agents. Firms can use tournaments to either hire or promote workers as well as to provide incentives for workers to exert effort. The seminal work on tournaments is Lazear and Rosen (1981).

Goel and Thakor (2008) study a two-period model of promotion tournaments where overconfident and rational managers compete. Overconfident managers underestimate the risk of their projects, i.e. overconfidence refers to overprecision. The principal does not know the managers' ability, but makes inferences based on the outcome of an investment project chosen by each manager. The manager with the highest project outcome becomes the new CEO. The main result is that an overconfident manager has a higher likelihood of being promoted to CEO than a rational one. This result hinges on how project risk maps onto outcomes: an overconfident manager inadvertently chooses a riskier project which means a lucky overconfident manager obtains a higher project outcome than a lucky rational manager. Since managerial ability, overconfidence, and project risk are unobservable, the principal cannot tell whether the better outcome is a consequence of higher ability or simply overconfidence, so her posterior belief about the ability of a lucky overconfident manager is higher than that of a lucky rational manager of the same ability. This implies that an overconfident manager is more likely to be promoted than a rational one.

Once a manager becomes the new CEO, the problem becomes one of information acquisition and project implementation, which is the topic of Section 4.1.6. The firm will prefer contracting with a moderately overconfident CEO than with a rational one, so a promotion rule that is biased towards overconfident CEOs, as was considered above, may well be optimal if the firm can either filter out or subsequently fire excessively overconfident CEOs.

Santos-Pinto (2010) studies how firms should design tournaments to provide incentives for overconfident employees to exert effort. The main finding is that firms can be better off with an overconfident workforce if they wisely structure tournament prizes.

As in Section 4.1.1, an overconfident worker overestimates the returns to effort. Worker overconfidence has two effects in tournaments used to incentivize employees' effort. First, it makes participation in a tournament more attractive to overconfident workers for a fixed prize structure.

An overconfident worker will overestimate the probability that he will attain the high prize, paid to the worker who produces the higher outcome. Second, it changes workers' incentives to exert effort for a fixed prize structure. Overconfidence raises effort provision when effort and self-confidence are complements but lowers it when effort and self-confidence are substitutes. These are the counterparts of the participation and incentive effects of overconfidence introduced in Section 4.1.1.

If workers are overconfident and effort and self-confidence are complements, the firm can implement the same effort with lower prizes or obtain a higher outcome for the same prizes. Hence, a firm's profits are higher in a tournament with an overconfident workforce than with a workforce with correct beliefs. If workers are overconfident and effort and self-confidence are substitutes, the firm needs to raise the power of incentives to implement the same effort level in a tournament with an overconfident workforce as in one with a workforce with correct beliefs. If workers are risk averse, they must be compensated for this increase in risk. If overconfidence is not too high, the firm is better off with an overconfident workforce. However, if overconfidence is high it can be detrimental to the firm because the firm must compensate the workers for the risk following an increase in the power of incentives.

Finally, overconfidence can lead to an overall welfare improvement. That is the case when overconfidence is moderate and workers have increasing absolute risk aversion. When effort is unobservable and workers have increasing absolute risk aversion and correct beliefs, there is under-supply of effort. If overconfidence lowers the prizes the firm needs to pay for workers' effort, the firm prefers to implement a higher effort level with a moderately overconfident workforce than with one with correct beliefs. This brings the second-best effort level closer to the first-best and therefore raises welfare.

#### 4.1.6 Investment Decisions with Private Information

Gervais et al. (2011) consider an employment relationship where a manager, instead of making an unobservable effort decision, decides whether the firm should pursue an investment project. The manager is risk averse and has private access to a signal which is informative about the probability of success or failure of the project; the signal is unobservable to the principal. Thus, instead of a hidden-action problem such as that in the basic model, Gervais et al. (2011) consider a hidden-information problem. The principal's incentive problem is to align the manager's investment decision with firm-value maximization. Managerial overconfidence (in this case about the precision of his private signal) makes it cheaper for the firm to solve this problem. Similar to the results in the basic model when effort and self-confidence are complements, a moderately overconfident manager will receive flatter compensation which will make him better off, while a significantly overconfident manager will receive more convex compensation, which will make him worse off. Gervais et al. (2011) extend the model, allowing for the private signal to be costly to the manager. The firm then needs to incentivize both information gathering and the correct investment decision given the observed signal. Since an overconfi-

dent manager overestimates the value of information in terms of decision making and his pay is contingent on the project outcome, it is cheaper for the firm to incentivize information gathering when contracting with an overconfident manager.

#### 4.1.7 Wage-Setting Policies

Fang and Moscarini (2005) analyze the impact of worker overconfidence on optimal wage-setting policies when the firm has private information about the skills of its workforce.

Agents differ in their skills which can be either low or high. For each agent, neither he nor the principal knows the true value of his skill. All agents have identical prior beliefs that they have high skill. The principal also holds identical prior beliefs about all the agents. Each agent is (weakly) overconfident about his skill relative to the principal's correct assessment. The prior beliefs are common knowledge, and when they are not equal, each party believes the other to be wrong.

Each employed agent can produce two levels of output, one normalized to zero and a positive amount. The production technology is stochastic: an increase in effort raises the probability of high output. The principal and the agents are risk neutral and the agents have limited liability. The effort of each agent is not observable to the principal, while his output is observable and verifiable by all parties. The principal, before beginning the relationship, receives a private performance evaluation (or signal) about each agent's skill. The principal uses the signal to update beliefs about each agent's skill. The principal decides, for each agent, whether to offer him a contract, and under which terms. A contract consists of a base wage, to be received regardless of the level of output, and a bonus, to be paid only when the output is positive. An agent can observe all contracts offered by the firm and compares his contract with others', makes inferences about the principal's performance evaluation, and adjusts his posterior beliefs accordingly. Note that since there is a continuum of agents, the Law of Large Numbers implies that the principal does not face any surprises at the aggregate level.

Fang and Moscarini (2005) assume the principal can only adopt one of two possible wage-setting policies: a non-differentiation policy, under which she offers the same wage contract to all agents irrespective of their performance evaluations, and a differentiation policy, under which different wage contracts are offered according to the agents' performance evaluations. The model generates two main results. First, in general the differentiation policy dominates the non-differentiation policy when there is no overconfidence and all parties share identical prior beliefs. Second, overconfidence can lead the principal to prefer a non-differentiation policy. Moreover, the more overconfident the workforce, and the lower its true average skill, the more likely the principal is to offer a non-differentiation policy. Intuitively, when agents are overconfident a differentiation policy conveys, on average, bad news to the workforce. Instead, a non-differentiation policy maintains a high level of overconfidence and effort provision by the workforce. This more than compensates for the foregone profits that could have been obtained, due to the sorting effect, under a

differentiation policy. This result might explain why many firms abstain from wage differentiation among their workers, as documented by Bewley (1999).

#### 4.1.8 Job Assignments

Nafziger (2011) asks how a firm should combine wage payments with job assignments when the firm can use the job assignment to signal whether an employee possesses the skills required to succeed in a job. A principal can employ an agent who can work in one of two jobs: job  $l$  and job  $r$ . In each job, the agent can either succeed or fail. If the agent succeeds, the principal obtains an observable and verifiable positive outcome. If the agent fails, the outcome is zero. The agent is risk-neutral, is protected by limited liability, and has a reservation utility of zero. The probability that the agent will succeed depends on the agent's effort and skill  $\theta \in [\theta_L, \theta_H]$ . Providing effort is costly and not providing it is costless. If the agent provides effort he will succeed at job  $j$  with probability  $p_j(\theta)$ . If the agent does not provide effort he will fail regardless of his skill. Interpreting  $p_j(\theta)$  as the agent's initial self-confidence, it follows that effort and self-confidence are complements. The higher the agent's skill, the higher the probability he succeeds at job  $r$ , and the lower the agent's skill, the higher the probability he succeeds at job  $l$ . These assumptions imply that there exists a unique cut-off  $\theta_E \in (0, 1)$  such that  $p_r(\theta_E) = p_l(\theta_E)$ . Hence, it is production-efficient to assign agents with skill lower than  $\theta_E$  to job  $l$ , and all others to job  $r$ .

At date 1, neither the agent nor the principal know the agent's skill. It is however common knowledge that skill is distributed according to a continuous density function which maps  $[\theta_L, \theta_H]$  to  $[0, 1]$ . The principal offers a contract to the agent. The contract specifies a wage if the revenue is high in job  $j$  and zero otherwise, and a job assignment rule which describes which of the two jobs the agent should work depending on his type. At date 2, the principal privately learns the agent's skill. At date 3, the principal assigns the agent to a job. The agent observes the job assignment and forms a posterior belief about his probability of success given the effort at job  $j$ . The posterior belief can be interpreted as the agent's interim self-confidence. At date 4, the agent provides unobservable effort. Note that as the principal specifies the wage before she learns the agent's skill, the wage cannot transmit any information from the principal to the agent. In contrast, the job assignment to a specific job may signal to the agent that the principal believes that he possesses the skills required for that job. As the job assignment rule influences the agent's interim self-confidence, the principal can use it to boost the agent's self-confidence and thus save on monetary rewards.

The separating equilibrium of this game is characterized by a unique cut-off  $\theta_S < \theta_E$  such that the principal assigns agents with a skill lower than  $\theta_S$  to job  $l$  and all others to job  $r$ . This implies that more agents are assigned to job  $r$ , the job where high-skill agents are more productive, than is efficient. Hence, the principal distorts the job-assignment rule to enhance interim self-confidence. The increase in interim self-confidence raises effort provision (due to the complementarity between effort and

self-confidence) and allows a reduction in the wage bill. This provides a rationale for why firms might not separate job assignments from the provision of incentives (see Baker et al., 1988).

#### 4.1.9 Subjective Performance Evaluation

The basic model shows how a firm can use an objective performance measure (the outcome of a project) to design a compensation scheme for an overconfident worker. However, in many jobs objective performance measures are extremely difficult to obtain (Prendergast, 1999; Zábojník, 2014). To provide incentives then, firms typically use subjective performance measures like evaluations of supervisors, co-workers, or consumers. The role of subjective performance evaluation (SPE) has been an intense subject of study in labor economics (MacLeod and Malcomson, 1989; Baker et al., 1994; Levin, 2003; MacLeod, 2003).

Foschi and Santos-Pinto (2018) show how a firm should design a SPE compensation scheme when a worker overestimates the precision of his private SPE signal. The main findings are as follows. First, the principal is better off with an agent who overestimates the precision of his SPE signal. Second, overprecision can either lead to exploitation or to a Pareto improvement. Third, when the agent's overprecision is significantly high, the optimal SPE compensation scheme uses the agent's SPE signal (or the agent's self-evaluation) as an input to set the agent's compensation. This goes against the general recommendation in the management literature to avoid the use of self-evaluations (or "self-assessments") in setting compensation (Milkovich et al., 2011; Bratton and Gold, 2012); a practice that is widely used nowadays.

## 4.2 Empirical Evidence

This section reviews empirical evidence on the impact of worker overconfidence on internal labor markets.

### 4.2.1 Experimental Evidence

Using field data to test how worker overconfidence affects labor markets is difficult since workers' effort, ability, and subjective beliefs about ability are generally not available. As a consequence, economists resort to laboratory experiments. The controlled environment of the laboratory allows economists to make causal claims about the impact of overconfidence on labor market decisions and outcomes. For a good summary of the advantages of laboratory experiments in comparison with data sets typically used in labor economics, see Falk and Gächter (2008).

Niederle and Vesterlund (2007) show that gender differences in overplacement can lead to gender differences in choice of compensation scheme. Despite there being no gender differences in performance, 73 percent of the men select to enter a tournament but only 35 percent of the women make this choice. The gender gap in tournament entry is driven by two factors. First, men have a stronger preference for competing than women.

Second, men are substantially more overconfident about their relative performance than women. About one third of the gender gap in tournament entry is accounted by gender differences in confidence. This experiment sheds light on the gender pay gap and why women are underrepresented in many high-profile jobs and across whole professions.

Dohmen and Falk (2011) study how overplacement and other personal characteristics lead participants to choose between a fixed and a variable compensation scheme (piece rate, tournament, or revenue sharing). The experiment finds that participants are more likely to select into tournaments the more they overplace themselves. Controlling for risk attitudes, men are not significantly more likely to select a variable payment.

An interesting result from the basic model is a non-trivial mapping from agent overconfidence to the power of incentives of the equilibrium compensation scheme. Given the possibility of both increasing and decreasing power of incentives with respect to overconfidence, the actual relationship is an interesting empirical question. Larkin and Leider (2012) conduct an experiment relating overestimation of absolute performance to the choice of compensation scheme. The experiment finds that highly-overconfident participants incorrectly choose a convex piece rate, over a linear one, more often than other participants. Larkin and Leider (2012) also find that participants who choose a convex piece rate do not perform differently from those randomly assigned to a convex piece rate, while those who choose a linear piece rate perform worse than those randomly assigned to a linear piece rate. These results suggest that firms can increase profits by attracting overconfident employees by offering convex compensation schemes.

Another interesting result from the basic model is that firms can exploit overconfident workers. Sautmann (2013) runs an experiment that finds support for this prediction. At stage 1, participants take a trivia quiz and are asked to guess their score. At the start of stage 2, participants are assigned to the role of principal or agent. Participants are ordered by level of self-confidence and the highest (lowest) quartile becomes the group of overconfident (underconfident) agents. The remaining participants become principals and are randomly assigned to either group. The main findings are as follows. First, the principal's expected profit increases when contracting with an overconfident agent, and decreases when hiring an underconfident agent, compared to an agent with correct beliefs. Second, all else equal, overconfident agents are offered—and accept—lower expected wages. In other words, principals are able to exploit overconfident agents. Third, on average, the self-confidence bias has no effect on the wage gap. Fourth, the few principals who raise the wage gap are able to successfully induce higher investment from overconfident agents.

#### 4.2.2 Field Evidence

Are workers overconfident about their productivity in an actual workplace setting? What are the consequences of this overconfidence in terms of workers' and firms' decisions? Evidence from CEOs, CFOs, and truck drivers shows overconfidence has consequences for real-world internal labor markets.

A strand of the literature focuses on the effects of overconfidence on corporate investment decisions. Malmendier and Tate (2005, 2008) classify overconfident CEOs as those who hold on to their stock options until the year of expiration, provided that the options are at least 40 percent in the money at the start of that year. Risk aversion and underdiversification predict that CEOs should exercise options immediately after the vesting period if the amount in the money is beyond a rational benchmark. CEOs who hold inflated views of their skills, however, overestimate the future performance of their firms and are therefore more willing to hold options, expecting to profit from future stock price increases. They also find that overconfident CEOs are 65 percent more likely to complete mergers, overpay for the target companies, and undertake value-destroying mergers. Malmendier and Tate (2015) update the measures of CEO overconfidence in Malmendier and Tate (2005, 2008) to more recent time periods, and find that the percentage of CEOs identified as overconfident lies around 40 percent. Ben-David et al. (2013) classify overconfident CFOs as those who display overly tight forecasts of own firms' cash flows. The study documents a positive correlation between this measure of CFO overconfidence and a measure based on the tightness of the same CFOs' estimates of the S&P 500 returns. The study also finds that firms with overconfident CFOs have 8 percent higher capital expenditures than the average firm and that overconfidence regarding long-term return distributions helps explain the level of capital expenditures and acquisitions.

One of the most consistent consequences of worker overconfidence in the theoretical models reviewed, is that one would expect principals to design compensation contracts that exploit such bias. Otto (2014) examines the effect of overconfidence on the level of CEO pay, finding a significant negative relationship between CEO overconfidence and the level of both total and incentive pay. This could indicate that overconfident CEOs get "the short end of the stick" more often than they anticipate, consistent with exploitation. Humphery-Jenner et al. (2016) use compensation data of CEOs between 1992 and 2011 to study the relationship between CEO overconfidence and the power of incentives as measured by option and equity intensity (the proportion of option and equity pay to total pay). There is evidence of increasing power of incentives in step with overconfidence. The wagering motive, which dominates for high levels of agent overconfidence, may thus be the driving force behind the design of compensation schemes for CEOs.

Hoffman and Burks (2019) use data on productivity and beliefs of truck drivers over two years, and show that truck drivers systematically and persistently overestimate their productivity, that overestimating productivity (in the current job) is tied to a lower likelihood to resign, and that overconfidence leads to higher profits for the firm and lower welfare for the employee. This field study provides additional support for the predictions of the basic model in Section 4.1.1.

## 5 External Labor Markets

This section reviews research on the impact of overconfidence on competitive labor markets.

### 5.1 Theory

There is a large literature on the gender pay gap—the fact that men earn, on average, more than women. Various explanations have been offered to justify its existence: discrimination against women (Becker, 1971; Rothschild and Stiglitz, 1982), gender differences in human capital (Mincer and Polachek, 1974), risk attitudes (Eckel and Grossman, 2008), preferences for competing, beliefs about relative performance (Niederle and Vesterlund, 2007), or women tending to cluster in lower-paid occupations, industries, and firms (Blau and Kahn, 2017).

Santos-Pinto (2012) shows that gender differences in self-confidence can lead to a gender pay gap in a competitive labor market. If education raises productivity, men underestimate the marginal cost of acquiring education, and women overestimate it, then men will, on average, earn more than women. Intuitively, before making any educational investments, men and women are equally productive. If men underestimate the marginal cost of acquiring education and women overestimate it, however, the proportion of men who acquire a high education level will be higher than the proportion of women who acquire a high education level. If education raises productivity, then men will be, on average, more productive than women, which generates the gender pay gap.

Biased beliefs about employment prospects can have an impact on job search and the optimal design of unemployment insurance. Spinnewijn (2015) reports a remarkable bias in a sample of U.S. unemployed job seekers' perceptions of job search. On average, unemployed job seekers expect to remain unemployed for an additional 6.8 weeks. However, the average remaining unemployment duration for the same sample of job seekers exceeded 23 weeks.

Spinnewijn (2015) studies the impact of biased beliefs on job search and unemployment insurance using a two-period model inspired by Baily (1978). A social planner decides the unemployment benefits and an employed risk-averse worker, taking these into account together with perceived employment prospects, decides in the first period how much to save to protect himself against the loss of earnings from unemployment. In the second period, the worker decides, if unemployed, how much (costly) effort to exert searching for a job. The worker overestimates the probability of finding a job either in the sense of *baseline optimism* (that is, an overestimation of the job finding probability for given search efforts) or *control optimism* (that is, an overestimation of the increase in the job finding probability when searching more intensively). In the first period, an optimistic worker underestimates the value of unemployment insurance and protects himself less against the unemployment risk through precautionary savings. In the second period, a baseline-optimistic worker searches too little and depletes his resources too fast, which has consequences for unemployment duration and implications for the design of unemployment

insurance. In contrast, a control-optimistic worker exerts more search effort in the second period and, as a consequence, his probability of finding work is higher. Hence, when the social planner chooses the level of unemployment benefits for a control-optimistic worker, she must account for the effect of an increase in search behavior on the worker's welfare, and respond to control optimism by increasing benefits to lower search effort.

## 5.2 Empirical Evidence

As Section 4.2 shows, laboratory experiments find that overplacement is an important factor for selection into competitive environments. However, it is less well understood how this result extrapolates to real-world settings. Buser et al. (2014), Wiswall and Zafar (2015), and Schulz and Thöni (2016) connect this finding to an important life decision which affects career paths and earnings: selecting an academic field of study.

Buser et al. (2014) test the predictive power of the Niederle and Vesterlund (2007) laboratory measure of competitiveness for the choice of academic track of a sample of secondary school students in the Netherlands. The study reports five main findings. First, although boys and girls display similar levels of academic ability, boys choose substantially more math- and science intensive academic tracks. Second, controlling for performance, girls are about 23 percentage points less likely to enter the tournament. Third, slightly over 30 percent of the gender gap in tournament entry can be explained by gender differences in confidence. Fourth, risk attitudes, whether measured by a lottery choice or a simple questionnaire item, significantly predict tournament entry but reduce the gender gap in tournament entry only by a small amount once confidence is controlled for. Fifth, the gender difference in tournament entry accounts for a substantial portion of the gender difference in academic track choice.

Wiswall and Zafar (2015) examine earnings perceptions (see Section 2.2.4) and choice of major of a sample of undergraduate college students at New York University. The study finds that while earnings are a significant determinant of major choice, beliefs about ability also matter. Women are considerably more likely than men to choose humanities/arts majors whereas men are more likely than woman to choose non-humanities majors. Ability differences between men and women explain part of the gender gap in major choice. Differences in beliefs about ability exacerbate the tendency for men to major in non-humanities subjects more than women.

Schulz and Thöni (2016) study the relation between self-placement and the choice of academic field of a sample of first year university students in Switzerland. In this study, participants take a trivia quiz where they have to guess the year of five historical events of the 20<sup>th</sup> century. Subsequently, participants rank themselves within their group in each session. Participants who estimate their rank correctly receive 5 CHF; any other estimate is not rewarded. On average, participants from Political Science, Law, Business Administration and Economics overplace themselves between 1.4 and 0.8 ranks, while participants from Humanities, Natural Science, or Medicine underplace themselves by 0.8, 0.3 and 0.2 ranks, respectively. The experiment also finds that a participant's self-placement

is a significant predictor of the choice of academic field.

An important empirical finding of the literature on job search is the negative duration dependence of exit rates out of unemployment. Understanding why employment prospects are worse for the long-term unemployed has generated an extensive literature in labor economics (Machin and Manning, 1999). The adverse welfare consequences of job loss are large, especially for those who suffer long spells of unemployment.

Mueller et al. (2018) examine job seekers' perceptions and how their relationship to unemployment outcomes contribute to the incidence of long-term unemployment in the U.S.. Job seekers' beliefs about their chances of re-employment are elicited at different time horizons using two longitudinal surveys. The study provides five main results. First, job seekers' stated beliefs have strong predictive power of actual job finding. Second, on average, job seekers overestimate their job finding chances. Third, job seekers report slightly higher job finding chances the longer they are unemployed, which runs counter to the empirical finding of negative duration dependence. Fourth, job seekers with a high (low) underlying job finding rate tend to underestimate (overestimate) their job finding chances. Fifth, a calibrated job search model shows that (i) correcting the biases in beliefs reduces the share of workers who are unemployed for longer than 6 months by 2-3 percentage points, and (ii) biases in beliefs jointly explain between 12 and 14 percent of the incidence of long-term unemployment.

## 6 Summary

This section summarizes the main findings of the chapter and discusses avenues for future research.

Section 2 reviews empirical research on overconfidence. The psychology evidence shows that, on the one hand, overconfidence depends on personal factors such as ability at a task and traits like extroversion or narcissism. On the other hand, overconfidence also depends on environmental factors such as specific characteristics of a task (e.g., how difficult it is) or the ambiguity of the ability under consideration. The economic evidence shows that overconfidence exists under incentives for truth telling, influences economic decisions, persists after repeated feedback, and can many times be irrational.

Section 3 surveys economic explanations for overconfidence. Rational overconfidence is compatible with Bayesian updating from a common prior. Irrational overconfidence can result from differences in prior beliefs or opinions. Overconfidence may also result from biases in information processing. In addition, overconfidence may provide benefits that compensate for its decision-making costs. For example, changing a rival's behavior to one's advantage, lowering moral costs of deceiving others, reducing self-control problems, increasing offspring, performing better at a stressful task, raising ego utility, or social approval. The diversity of economic explanations for overconfidence, makes it less surprising why this is such an ubiquitous phenomenon.

Section 4 reviews research on the impact of worker overconfidence on

internal labor markets. A basic principal-agent model shows how firms can design compensation schemes to take advantage of worker overconfidence. Hence, the basic model shows overconfident workers can be more desirable employees than workers with either correct or underconfident beliefs. This implies that firms will treat overconfident workers differently leading to systematic differences in labor market outcomes among groups which vary in self-confidence. This result generalizes to employment relationships like teamwork, tournaments, investment decisions, and subjective performance evaluations. In addition, worker overconfidence can often lead to Pareto improvements when multiple firms compete for a single employee, employees are risk averse, or when there are complementarities in production. Section 4 also shows that the main predictions of the basic model are supported by experimental evidence. Furthermore, field evidence from CEOs, CFOs, and truck drivers shows worker overconfidence has consequences for real-world internal labor markets.

Section 5 reviews research on how worker overconfidence affects external labor markets. Gender differences in self-confidence can lead to a gender pay gap in competitive labor markets. Worker overconfidence and optimism affect job search and the design of optimal unemployment insurance. Empirical studies show worker overconfidence affects external labor markets. Self-confidence plays a role in the selection of an academic field of study and gender differences in self-confidence are correlated with gender differences in academic track choice. Additionally, overestimation of the probability of finding a job contributes to the slow exit out of unemployment and the incidence of long-term unemployment. These empirical findings have important consequences for career paths and earnings.

Understanding how worker overconfidence affects labor markets is a new and exciting area of research. This literature has been growing steadily, and picking up notoriety in the last decade. There are several promising avenues for future research on this topic. On the theory side, the interaction between worker overconfidence and private information has not yet been fully analyzed. There is a need for studies that offer a systematic treatment of the three different forms of overconfidence: overestimation, overplacement, and overprecision. The theory also has to expand beyond the unobservable effort framework in order to describe the effects of worker overconfidence on a broader set of labor market relationships.

On the empirical side, there is a need for further tests regarding the impact of worker overconfidence on effort provision and the design of compensation schemes. The basic model predicts that the power of incentives of the optimal compensation scheme is determined jointly by a worker's degree of overconfidence and by the relation between effort and self-confidence (whether these are complements or substitutes). This makes for rich but challenging possibilities of "testing the model", which the empirical and experimental literature has only recently started to exploit. More research about worker overconfidence in the field is also needed, especially in the context of employee productivity in actual workplace settings, job search, human capital investment, and on how worker overconfidence affects firms' hiring, promotion, and wage policies.

## Cross References

Behavioral job search  
Bonus plans, subjective performance evaluation and career concerns  
Tournaments and competition  
Wage policies, incentive schemes and motivation  
Feedback policies and peer effects at work  
Gender in the labor market

## Acknowledgments

We thank Adrian Bruhin, José Mata, Joel Sobel, the editor, Marie Claire Villeval, and an anonymous referee for helpful comments and suggestions.

## References

- Alicke, M. D. (1985). Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology*, 49(6):1621–1630.
- Alicke, M. D. and Govorun, O. (2005). The better-than-average effect. *The Self in Social Judgment*, 1:85–106.
- Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402.
- Baker, G., Gibbons, R., and Murphy, K. J. (1994). Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, 109(4):1125–1156.
- Baker, G. P., Jensen, M. C., and Murphy, K. J. (1988). Compensation and incentives: Practice vs. theory. *The Journal of Finance*, 43(3):593–616.
- Banerjee, R., Gupta, N. D., and Villeval, M. C. (2019). Self confidence spillovers, status and motivated beliefs. GATE Working Paper 1806.
- Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232.
- Becker, G. S. (1971). *The Economics of Discrimination*. Chicago University Economics Research Center. Economics Research Studies. Chicago University Press, Chicago, IL.
- Ben-David, I., Graham, J. R., and Harvey, C. R. (2013). Managerial miscalibration. *The Quarterly Journal of Economics*, 128(4):1547–1584.
- Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915.
- Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678.

- Benoit, J.-P. and Dubra, J. (2011). Apparent overconfidence. *Econometrica*, 79(5):1591–1625.
- Benoit, J.-P., Dubra, J., and Moore, D. A. (2015). Does the better-than-average effect show that people are overconfident?: Two experiments. *Journal of the European Economic Association*, 13(2):293–329.
- Bewley, T. (1999). *Why don't wages fall in a recession*. Harvard University Press, Cambridge, MA.
- Biais, B., Hilton, D., Mazurier, K., and Pouget, S. (2005). Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. *The Review of Economic Studies*, 72(2):287–312.
- Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865.
- Borch, K. (1962). Equilibrium in a reinsurance market. *Econometrica*, 30(3):424–444.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2019). Beliefs about gender. *American Economic Review*, 109(3):739–773.
- Bratton, J. and Gold, J. (2012). *Human resource management: theory and practice*. Palgrave Macmillan, New York, NY.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3.
- Bruhin, A., Fehr-Duda, H., and Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412.
- Bruhin, A., Manai, M., and Santos-Pinto, L. (2019). Risk and rationality: The relative importance of probability weighting and choice set dependence. Working Paper, University of Lausanne.
- Bruhin, A., Santos-Pinto, L., and Staubli, D. (2018). How do beliefs about skill affect risky decisions? *Journal of Economic Behavior & Organization*, 150:350–371.
- Burks, S. V., Carpenter, J. P., Goette, L., and Rustichini, A. (2013). Overconfidence and social signalling. *The Review of Economic Studies*, 80(3):949–983.
- Buser, T., Gerhards, L., and Van Der Weele, J. (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty*, 56(2):165–192.
- Buser, T., Niederle, M., and Oosterbeek, H. (2014). Gender, competitiveness, and career choices. *The Quarterly Journal of Economics*, 129(3):1409–1447.

- Cain, D. M., Moore, D. A., and Haran, U. (2015). Making sense of overconfidence in market entry. *Strategic Management Journal*, 36(1):1–18.
- Camerer, C. and Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1):306–318.
- Campbell, W. K., Goodie, A. S., and Foster, J. D. (2004). Narcissism, confidence, and risk attitude. *Journal of Behavioral Decision Making*, 17(4):297–311.
- Carrillo, J. D. and Mariotti, T. (2000). Strategic ignorance as a self-disciplining device. *The Review of Economic Studies*, 67(3):529–544.
- Charness, G., Rustichini, A., and Van de Ven, J. (2018). Self-confidence and strategic behavior. *Experimental Economics*, 21(1):72–98.
- Chew, S. H., Huang, W., and Zhao, X. (2018). Motivated false memory. Working Paper, National University of Singapore.
- Christensen-Szalanski, J. J. and Bushyhead, J. B. (1981). Physicians’ use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human Perception and Performance*, 7(4):928–935.
- Clark, J. and Friesen, L. (2009). Overconfidence in forecasts of own performance: An experimental study. *The Economic Journal*, 119(534):229–251.
- Compte, O. and Postlewaite, A. (2004). Confidence-enhanced performance. *American Economic Review*, 94(5):1536–1557.
- Conte, A., Hey, J. D., and Moffatt, P. G. (2011). Mixture models of choice under risk. *Journal of Econometrics*, 162(1):79–88.
- Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.
- Cross, K. P. (1977). Not can, but will college teaching be improved? *New Directions for Higher Education*, 1977(17):1–15.
- De Bondt, W. F. M. and Thaler, R. H. (1995). Financial decision-making in markets and firms: A behavioral perspective. In *Finance*, volume 9 of *Handbooks in Operations Research and Management Science*, chapter 13, pages 385–410. Elsevier.
- de la Rosa, L. E. (2008). Overconfidence and moral hazard. Working Paper, University of Aarhus.
- de la Rosa, L. E. (2011). Overconfidence and moral hazard. *Games and Economic Behavior*, 73(2):429–451.
- Deaves, R., Lüders, E., and Luo, G. Y. (2009). An experimental test of the impact of overconfidence and gender on trading activity. *Review of Finance*, 13(3):555–575.

- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372.
- DellaVigna, S. and Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review*, 96(3):694–719.
- Dohmen, T. and Falk, A. (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. *American Economic Review*, 101(2):556–590.
- Dunning, D. (1993). Words to live by: The self and definitions of social concepts and categories. In *Psychological perspectives on the self*, volume 4. The self in social perspective, pages 99–126. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Dunning, D. and Hayes, A. F. (1996). Evidence for egocentric comparison in social judgment. *Journal of Personality and Social Psychology*, 71(2):213–229.
- Dunning, D., Meyerowitz, J. A., and Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57(6):1082–1090.
- Dunning, D., Perie, M., and Story, A. L. (1991). Self-serving prototypes of social categories. *Journal of Personality and Social Psychology*, 61(6):957–968.
- Eckel, C. C. and Grossman, P. J. (2008). Men, women and risk aversion: Experimental evidence. In Plott, C. R. and Smith, V. L., editors, *Handbook of Experimental Economics Results*, volume 1, chapter 113, pages 1061–1073. Elsevier.
- Ederer, F. and Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science*, 59(7):1496–1513.
- Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–38.
- Erev, I., Wallsten, T. S., and Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3):519–527.
- Ewers, M. and Zimmermann, F. (2015). Image and misreporting. *Journal of the European Economic Association*, 13(2):363–380.
- Falk, A. and Gächter, S. (2008). Experimental labour economics. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*, pages 1–5. Palgrave Macmillan, London, UK.
- Fang, H. and Moscarini, G. (2005). Morale hazard. *Journal of Monetary Economics*, 52(4):749–777.

- Felson, R. B. (1981). Ambiguity and bias in the self-concept. *Social Psychology Quarterly*, pages 64–69.
- Fischhoff, B., Slovic, P., and Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 3(4):552–564.
- Foschi, M. and Santos-Pinto, L. P. (2018). Subjective performance evaluation of employees with biased beliefs. Working Paper, European University Institute.
- Gervais, S. and Goldstein, I. (2007). The positive effects of biased self-perceptions in firms. *Review of Finance*, 11(3):453–496.
- Gervais, S., Heaton, J. B., and Odean, T. (2011). Overconfidence, compensation contracts, and capital budgeting. *The Journal of Finance*, 66(5):1735–1777.
- Gervais, S. and Odean, T. (2001). Learning to be overconfident. *The Review of Financial Studies*, 14(1):1–27.
- Gigerenzer, G., Hoffrage, U., and Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98(4):506–528.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1):384–394.
- Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, 108(2):419–53.
- Goel, A. M. and Thakor, A. V. (2008). Overconfidence, CEO selection, and corporate governance. *The Journal of Finance*, 63(6):2737–2784.
- Griffin, D. and Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3):411–435.
- Grossman, Z. and Owens, D. (2012). An unlucky feeling: Overconfidence and noisy feedback. *Journal of Economic Behavior & Organization*, 84(2):510–524.
- Grubb, M. D. (2009). Selling to overconfident consumers. *American Economic Review*, 99(5):1770–1807.
- Gul, F. (1998). A comment on Aumann’s Bayesian view. *Econometrica*, 66(4):923–927.
- Guthrie, C., Rachlinski, J. J., and Wistrich, A. J. (2001). Inside the judicial mind. *Cornell Law Review*, 86:777–830.
- Harless, D. W. and Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, pages 1251–1289.

- Haun, D. E., Zeringue, A., Leach, A., and Foley, A. (2000). Assessing the competence of specimen-processing personnel. *Laboratory Medicine*, 31(11):633–637.
- Heath, C. and Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1):5–28.
- Heifetz, A., Shannon, C., and Spiegel, Y. (2007a). The dynamic evolution of preferences. *Economic Theory*, 32(2):251–286.
- Heifetz, A., Shannon, C., and Spiegel, Y. (2007b). What to maximize if you must. *Journal of Economic Theory*, 133(1):31–57.
- Herz, H., Schunk, D., and Zehnder, C. (2014). How do judgmental overconfidence and overoptimism shape innovative activity? *Games and Economic Behavior*, 83:1–23.
- Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, pages 1291–1326.
- Hodges, B., Regehr, G., and Martin, D. (2001). Difficulties in recognizing one’s own incompetence: novice physicians who are unskilled and unaware of it. *Academic Medicine*, 76(10):S87–S89.
- Hoelzl, E. and Rustichini, A. (2005). Overconfident: Do you put your money on it? *The Economic Journal*, 115(503):305–318.
- Hoffman, M. and Burks, S. V. (2019). Worker overconfidence: Field evidence and implications for employee turnover and firm profits. Forthcoming, *Quantitative Economics*.
- Hoffrage, U. (2016). Overconfidence. In Pohl, R. F., editor, *Cognitive illusions: Intriguing phenomena in judgement, thinking and memory*, chapter 13. Psychology Press, New York, NY.
- Holmstrom, B. (1982). Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340.
- Hossain, T. and Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, 80(3):984–1001.
- Huang, L. and Murad, Z. (2017). Impact of relative performance feedback on beliefs and preferences across dissimilar tasks. Better for Less Discussion Paper Series No. 02/2016, University of Surrey.
- Huffman, D., Raymond, C., and Shvets, J. (2019). Persistent overconfidence and biased memory: Evidence from managers. Working paper, University of Pittsburgh.
- Humphery-Jenner, M., Lisic, L. L., Nanda, V., and Silveri, S. D. (2016). Executive overconfidence and compensation structure. *Journal of Financial Economics*, 119(3):533–558.

- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2):603–606.
- Kőszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707.
- Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature*, 52(4):1075–1118.
- Klayman, J. and Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94(2):211–228.
- Koriat, A., Lichtenstein, S., and Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2):107–118.
- Krueger, J. and Mueller, R. A. (2002). Unskilled, unaware, or both? the better-than-average heuristic and statistical regression predict errors in estimates of own performance. *Journal of Personality and Social Psychology*, 82(2):180–188.
- Kruger, J. (1999). Lake wobegon be gone! the “below-average effect” and the egocentric nature of comparative ability judgments. *Journal of Personality and Social Psychology*, 77(2):221–232.
- Kruger, J. and Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121–1134.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3):480–498.
- Laffont, J. and Martimort, D. (2002). *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton, NJ.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.
- Larkin, I. and Leider, S. (2012). Incentive schemes, sorting, and behavioral biases of employees: Experimental evidence. *American Economic Journal: Microeconomics*, 4(2):184–214.
- Larwood, L. and Whittaker, W. (1977). Managerial myopia: Self-serving biases in organizational planning. *Journal of Applied Psychology*, 62(2):194–198.
- Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864.
- Levin, J. (2003). Relational incentive contracts. *American Economic Review*, 93(3):835–857.

- Lichtenstein, S., Fischhoff, B., and Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In Kahneman, D., Slovic, P., and Tversky, A., editors, *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11):2098–2109.
- Machin, S. and Manning, A. (1999). The causes and consequences of longterm unemployment in Europe. In *Handbook of Labor Economics*, volume 3, chapter 47, pages 3085–3139. Elsevier.
- MacLeod, W. B. (2003). Optimal contracting with subjective evaluation. *American Economic Review*, 93(1):216–240.
- MacLeod, W. B. and Malcomson, J. M. (1989). Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica*, 57(2):447–480.
- Malmendier, U. and Tate, G. (2005). CEO overconfidence and corporate investment. *The Journal of Finance*, 60(6):2661–2700.
- Malmendier, U. and Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market’s reaction. *Journal of Financial Economics*, 89(1):20–43.
- Malmendier, U. and Tate, G. (2015). Behavioral CEOs: The role of managerial overconfidence. *Journal of Economic Perspectives*, 29(4):37–60.
- Malmendier, U., Tate, G., and Yan, J. (2011). Overconfidence and early-life experiences: The effect of managerial traits on corporate financial policies. *The Journal of Finance*, 66(5):1687–1733.
- Merkle, C. and Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Decision Processes*, 116(2):262–271.
- Milkovich, G., Newman, J., and Gerhart, B. (2011). *Compensation*. McGraw-Hill Education, New York, NY.
- Miller, D. T. and Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction? *Psychological bulletin*, 82(2):213–222.
- Mincer, J. and Polachek, S. (1974). Family investments in human capital: Earnings of women. *Journal of Political Economy*, 82(2, Part 2):S76–S108.
- Mobius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. Working Paper 17014, National Bureau of Economic Research.

- Moore, D. A. and Cain, D. M. (2007). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes*, 103(2):197–213.
- Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2):502–517.
- Moore, D. A. and Kim, T. G. (2003). Myopic social prediction and the solo comparison effect. *Journal of Personality and Social Psychology*, 85(6):1121–1135.
- Moore, D. A., Tenney, E. R., and Haran, U. (2015). Overprecision in judgment. In Keren, G. and Wu, G., editors, *The Wiley Blackwell Handbook of Judgment and Decision Making*, volume II, chapter 6, pages 182–209. John Wiley & Sons, Ltd.
- Morris, S. (1995). The common prior assumption in economic theory. *Economics & Philosophy*, 11(2):227–253.
- Mueller, A. I., Spinnewijn, J., and Topa, G. (2018). Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence and bias. Working Paper 25294, National Bureau of Economic Research.
- Murad, Z., Sefton, M., and Starmer, C. (2016). How do risk attitudes affect measured confidence? *Journal of Risk and Uncertainty*, 52(1):21–46.
- Myers, D. G. (1998). *Social psychology*. McGraw-Hill.
- Nafziger, J. (2011). Motivational job assignments. *Economica*, 78(312):676–696.
- Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *The Quarterly Journal of Economics*, 122(3):1067–1101.
- Nisbett, R. E. and Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Prentice-Hall, Englewood Cliffs, NJ.
- O’Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.
- Oskamp, S. (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology*, 29(3):261–265.
- Otto, C. A. (2014). CEO optimism and incentive compensation. *Journal of Financial Economics*, 114(2):366–404.
- Park, Y. J. and Santos-Pinto, L. (2010). Overconfidence in tournaments: evidence from the field. *Theory and Decision*, 69(1):143–166.

- Phelps, E. S. and Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2):185–199.
- Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature*, 37(1):7–63.
- Preston, C. E. and Harris, S. (1965). Psychology of drivers in traffic accidents. *Journal of Applied Psychology*, 49(4):284–288.
- Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82.
- Ring, P., Neyse, L., David-Barett, T., and Schmidt, U. (2016). Gender differences in performance predictions: Evidence from the cognitive reflection test. *Frontiers in Psychology*, 7:1680. Article 1680.
- Rothschild, M. and Stiglitz, J. E. (1982). A model of employment outcomes illustrating the effect of the structure of information on the level and distribution of income. *Economics Letters*, 10(3):231–236.
- Russo, J. E. and Schoemaker, P. J. H. (1992). Managing overconfidence. *Sloan Management Review*, 33(2):7–17.
- Sanchez, C. and Dunning, D. (2018). Overconfidence among beginners: Is a little learning a dangerous thing? *Journal of Personality and Social Psychology*, 114(1):10–28.
- Santos-Pinto, L. (2008). Positive self-image and incentives in organisations. *The Economic Journal*, 118(531):1315–1332.
- Santos-Pinto, L. (2010). Positive self-image in tournaments. *International Economic Review*, 51(2):475–496.
- Santos-Pinto, L. (2012). Labor market signaling and self-confidence: Wage compression and the gender pay gap. *Journal of Labor Economics*, 30(4):873–914.
- Santos-Pinto, L. and Sobel, J. (2005). A model of positive self-image in subjective assessments. *American Economic Review*, 95(5):1386–1402.
- Sautmann, A. (2013). Contracts for agents with biased beliefs: Some theory and an experiment. *American Economic Journal: Microeconomics*, 5(3):124–156.
- Schaefer, P. S., Williams, C. C., Goodie, A. S., and Campbell, W. K. (2004). Overconfidence and the big five. *Journal of Research in Personality*, 38(5):473–480.
- Schlag, K. H., Tremewan, J., and Van der Weele, J. J. (2015). A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics*, 18(3):457–490.

- Schlag, K. H. and van der Weele, J. J. (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, 3(1):38–42.
- Schlag, K. H. and van der Weele, J. J. (2015). A method to elicit beliefs as most likely intervals. *Judgment and Decision Making*, 10(5):456–468.
- Schulz, J. F. and Thöni, C. (2016). Overconfidence and career choice. *PLoS ONE*, 11(1):e0145126.
- Schwardmann, P. and Van der Weele, J. J. (2019). Deception and self-deception. *Nature Human Behaviour*, 3:1055–1061.
- Solda, A., Ke, C., Page, L., and von Hippel, B. (2019). Strategically delusional. GATE Lyon Saint-Etienne, Working Paper 1908.
- Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1):130–167.
- Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38(2):332–382.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180.
- Svenson, O. (1981). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica*, 47(2):143–148.
- Taylor, S. E. and Brown, J. D. (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin*, 103(2):193–210.
- Trivers, R. (1985). *Social evolution*, chapter 16: Deceit and Self-Deception, pages 395–420. Benjamin-Cummings, Menlo Park, CA.
- Van den Steen, E. (2004). Rational overoptimism (and other biases). *American Economic Review*, 94(4):1141–1151.
- Van den Steen, E. (2011). Overconfidence by bayesian-rational agents. *Management Science*, 57(5):884–896.
- Von Hippel, W. and Trivers, R. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences*, 34(1):1–16.
- Waldman, M. (1994). Systematic errors and the theory of natural selection. *American Economic Review*, 84(3):482–497.
- Wiswall, M. and Zafar, B. (2015). Determinants of college major choice: Identification using an information experiment. *The Review of Economic Studies*, 82(2):791–824.
- Yates, J. F. (1990). *Judgment and decision making*. Prentice-Hall, Inc.

- Zábojník, J. (2004). A model of rational bias in self-assessments. *Economic Theory*, 23(2):259–282.
- Zábojník, J. (2014). Subjective evaluations with performance feedback. *The RAND Journal of Economics*, 45(2):341–369.
- Zenger, T. R. (1992). Why do employers only reward extreme performance? Examining the relationships among performance, pay, and turnover. *Administrative Science Quarterly*, 37(2):198–219.
- Zimmermann, F. (2019). The dynamics of motivated beliefs. Forthcoming, *American Economic Review*.