Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Firm Profits[*]

Mitchell Hoffman[†]
Stephen V. Burks[‡]

March 2019

Abstract

Combining weekly productivity data with weekly productivity beliefs for a large sample of truckers over two years, we show that workers tend to systematically and persistently over-predict their productivity. If workers are overconfident about their own productivity at the current firm relative to their outside option, they should be less likely to quit. Empirically, all else equal, having higher productivity beliefs is associated with an employee being less likely to quit. To study the implications of overconfidence for worker welfare and firm profits, we estimate a structural learning model with biased beliefs that accounts for many key features of the data. While worker overconfidence moderately decreases worker welfare, it also substantially increases firm profits.

JEL Classifications: J24, D03, M53, J41
Keywords: Overconfidence; biased learning; turnover

1 Introduction

Scholars have long argued that people have a tendency to be overconfident about their ability (see, e.g., Adam Smith's *The Wealth of Nations*). Decades of research in psychology (and growing research in economics) support the idea that people are overconfident. However, much of this research is based on short-term student lab experiments. Much less is known about overconfidence in the field (especially over time) and even less so in the context of employee productivity in the workplace.[1] Are workers overconfident about their productivity in an actual workplace setting? Is overconfidence persistent or does it quickly disappear over time due to learning? What are the implications of worker overconfidence for employee behavior, employee welfare, and the profits of firms?

We address these questions using unique data from the trucking industry. While it is one industry, trucking is well-suited for our analysis because productivity (weekly miles driven) is easily measured and because the industry is large (see Section 2). At a leading trucking firm (which we call Firm A), 895 newly trained workers were asked to predict their weekly productivity for two years. We show that workers who expect higher productivity end up achieving higher productivity, so subjective beliefs are predictive. However, the data also reveal a pattern where workers tend to systematically overpredict their productivity. Overprediction is very persistent. Average overprediction does eventually decline, but only very slowly. The overprediction we observe without financial incentives remains even when belief-elicitation is made incentive-compatible using randomized financial incentives for accurate prediction at a second large trucking firm (Firm B). We refer to this overprediction as "overconfidence" and say more about the term below.

Having documented this overprediction, we next seek to model it quantitatively, as well as to understand its implications. We turn to Jovanovic's (1979) canonical model of turnover, where quitting decisions reflect the evolution of worker beliefs about job match or productivity. We document

---

[1]Some exceptions in economics on overconfidence in the field include the work on overconfident CEOs pioneered by Malmendier and Tate (2005), as well as Hoffman (2016), who studies how overconfidence affects businesspeople's demand for information; Wang (2015), who studies loan officers, accommodating biased beliefs in screening ability using a structural model; and work on overconfident investors (e.g., Barber and Odean, 2001). In psychology, there are various studies that examine overconfidence among particular workers, e.g., Baumann et al. (1991) study doctors and nurses. However, studies like Baumann et al. (1991) often measure beliefs only once and often consider hypothetical situations/vignettes or trivia questions instead of predicting productivity. Meikle et al. (2016) review work on overconfidence and organizations. To our knowledge, very few studies analyze overconfidence at high frequencies over substantial periods of time (e.g., > 3 months). An exception is a psychology study by Massey et al. (2011), who show that football fans persistently over-predict the chance of their favorite team winning over 4 months. Other papers study forms of biased beliefs over long time horizons, but at lower frequencies, e.g., quarterly (Ben-David et al., 2013).

that, consistent with theory, workers who expect higher productivity are less likely to quit. From the standpoint of the firm, this may be especially important in our context because the firm is providing the workers with firm-sponsored general training at no direct cost. Turnover is costly for the firm, leading the firm to lose the individuals that they recently provided training to. While potentially useful for firms, if workers are overconfident about their ability at the firm relative to their outside option, this may distort worker quitting decisions, reducing worker welfare.

To evaluate the importance of overconfidence for worker welfare and firm profits, we develop a structural model of worker turnover. Similar to Jovanovic (1979), workers learn about their underlying productivity through weekly productivity realizations, and decide when, if ever, to quit. However, we do not impose that workers are fully rational. Workers may hold biased priors, or learn faster or slower than predicted by Bayes' rule, nesting the standard model as a special case. Using our rich subjective belief data for identification, we estimate that workers have mean bias of about one-third of underlying productivity, as well as substantial variance bias, with learning much slower than predicted by Bayes' rule. Our model fits the data quite well, whereas a standard model performs far worse. In a counterfactual simulation, we show that eliminating worker overconfidence would moderately increase worker welfare (because workers make better decisions), but would substantially reduce firm profits.

Our study makes three main contributions to the literature. First, we provide long-term high-frequency field evidence on overconfidence, some of the longest high-frequency evidence in any field (psychology or economics).[2] Moore and Healy (2008) provide an excellent survey of recent work and divide overconfidence into three types: relative overconfidence (thinking you are better than others), absolute overconfidence (thinking you are better than you actually are), and over-precision (thinking your beliefs are more precise than they actually are). Our paper's largest focus is on absolute overconfidence, which we refer to hereafter simply as overconfidence. Much of overconfidence research studies short-term laboratory tasks, e.g., trivia games. This paper analyzes overconfidence using weekly data over two years on forecasts about individual productivity in an actual work setting.

Second, we quantify the worker welfare impacts of overconfidence by developing a structural

---

[2]To our knowledge, our study provided the longest high-frequency field evidence in the literature when it first appeared. In recent work by psychologists, Moore et al. (2016) study a geopolitical forecasting tournament, where people participated for up to 3 years, building on earlier work on political forecasting (Tetlock, 2005). Moore et al. (2016) find a small but persistent degree of overconfidence. Our study differs in that it examines workplace productivity (instead of geopolitics), it studies implications of overconfidence, and it models overconfidence using a structural model.

learning model with biased beliefs. We present one of the first papers in economics to estimate a learning model with biased beliefs.[3] More generally, we contribute to a small but growing literature using subjective beliefs in various ways to estimate structural models (for pioneer papers, see, e.g., Bellemare et al., 2008; Chan et al., 2008; van der Klaauw and Wolpin, 2008).[4]

Third, we show that worker overconfidence can significantly enhance firm profits. Counter-factual simulations suggest that biased beliefs are quantitatively important for firms; in particular, training would be substantially less profitable if workers were not overconfident. While a number of field studies analyze how firms may benefit from consumer biases (see Koszegi (2014) for a survey), ours is one of the first to analyze how firms may benefit from biases of their workers.[5]

The paper proceeds as follows. Section 2 gives background on trucking and describes the data. Section 3 analyzes subjective productivity belief data, both from Firm A (without incentives) and from Firm B (with randomized financial incentives). Section 4 develops the model and structurally estimates it. Section 5 performs the counterfactual simulations. Section 6 concludes.

2 Background and Data

2.1 Institutional Background

Truckdriving in the US. Truckdriving is a large occupation, with roughly 1.8 million US workers operating heavy trucks such as those used by the firms we study (BLS, 2010). Firms A and B are in the long-distance truckload segment of the for-hire trucking industry, which is the largest employment setting for this occupation. An important distinction is between long-haul and short-haul trucking. Long-haul truckload drivers are usually paid by the mile (a piece rate) (Belzer, 2000) and drive long distances from home. In contrast, short-haul truckload drivers generally spend fewer nights away

---

[3]While several recent papers in labor and personnel economics analyze learning using a structural approach (e.g., Arcidiacono, 2004; Bojilov, 2017; Sanders, 2016; Stange, 2012), we allow for both generalized and non-rational learning. Two papers in industrial organization, Goettler and Clay (2011) and Grubb and Osborne (2015), estimate biased learning models of plan choice for online groceries and cell phone service, respectively. A main difference in our paper is that belief biases are identified using high-frequency subjective belief data, whereas in Goettler and Clay (2011) and Grubb and Osborne (2015), biases are identified through contractual choices. There are advantages of each approach. An advantage of using contracts relative to using subjective beliefs is that economists are more trusting of "what people do" compared to "what people say." A virtue of using beliefs is that repeated sub-optimal *ex post* contractual choices may reflect factors other than biased beliefs, including inertia or switching costs.

[4]See Arcidiacono et al. (2014) and Wiswall and Zafar (2015) for examples of recent papers. van der Klaauw (2012) discusses incorporating subjective beliefs into dynamic structural models. Appendix A.13 describes additional papers.

[5]Concurrent with or after our paper first appeared, Otto (2014) and Humphery-Jenner et al. (2016) show that overconfident executives receive more performance pay.

from home and are not usually paid by the mile.[6]

The main training for heavy truckdrivers is that needed to obtain a commercial driver's license (CDL). Most new drivers take a formal CDL training course, and in some states it is required by law (BLS, 2010). CDL training can be obtained at truck driving schools run by trucking companies, at private truck driving schools, and at some community colleges. At Firm A, the CDL training drivers received went for about 2-3 weeks, and included classroom lectures, simulator driving, and actual behind-the-wheel truck driving. The market price for CDL training at private training schools varies, but is often several thousand dollars.

The drivers we study in this paper received training under a 12-month training contract. Under this contract, Firm A paid for the training and in return the driver committed to stay with the firm for a year. If the driver left early, they were fined between $3,500 and $4,000. Drivers did not post a bond, and the firm seemed to collect only about 30% of the penalties owed (despite the firm making strenuous efforts to collect the penalties owed); further details on the contracts are provided in another paper (Hoffman and Burks, 2017), which studies the contracts in detail.[7]

Production. Truckload drivers haul full loads between a wide variety of locations. While our data do not contain driver hours, drivers are constrained by the federal legal limit of about 60 hrs/week, and managers informed us that drivers often work up to the limit. Firm A loads are assigned via a central dispatching system and are assigned primarily by proximity (as well as hours left up to the federal limit). Once a load is finished, a driver may start a new one.

Productivity in long-haul trucking is measured in miles per week. There are significant cross-driver differences in average productivity, as well as substantial idiosyncratic variation in productivity within drivers. When asked the reason for significant cross-driver differences, managers described various factors including speed, skill at avoiding traffic, route planning (miles are calculated according to a pre-specified distance between two points, not by distance traveled), not getting lost, and coordinating with people to unload the truck. As an example, if a driver arrives late, they may

---

[6]We highlight a few more institutional details. Truckload is the segment that hauls full trailer loads. Truckload has employee turnover rates, often over 100% per year (Burks et al., 2008), as well as low unionization, and most drivers do not own their own trucks. Around 10% of trucks in 1992 were driven by drivers who own their own truck (owner-operators), with the remainder driven by drivers driving company-owned trucks (company drivers) (Baker and Hubbard, 2004). All the drivers we study are non-union company drivers. For an analysis of productivity in trucking, see Hubbard (2003).

[7]The Appendix of Hoffman and Burks (2017) explains how it is common for large truckload firms to provide CDL training.

need to wait around in order to have their truck unloaded, and this could negatively affect miles per week. As for the sources of week-to-week variation, managers emphasized weather, traffic, variable loading/unloading time, and disadvantageous load assignments. Thus, weekly miles, our measure of productivity, reflect both driver performance and effort, as well as factors that drivers do not control and may be difficult to predict *ex ante*. See Appendix G for more on measuring productivity.

2.2 Firm A Data

Data Information. To create our dataset, we collected subjective beliefs about next week's productivity for a subset of 895 new drivers trained at one of the firm's training schools in late 2005-2006. Beyond the productivity beliefs survey, drivers did various tests (e.g., IQ, personality) during training, and were invited to do other surveys during their first two years of work (see Appendix A.1). We will sometimes refer to drivers in our data as the "data subset,"[8] and several other papers by the author analyze this subset of drivers in other work.[9] However, the productivity belief data we collect have never been analyzed previously. Records from the firm provide weekly data on miles and earnings, and we also have worker demographic information.

Every week around Tuesday,[10] drivers in the data subset were asked to predict their miles for the following pay week (Sunday-Saturday, starting on the Sunday in 5 days). This occurred for up to roughly two years, with some variation in maximum weeks depending on when drivers started. Drivers responded by typing an answer to the below question, which we sent over the truck's computer system: *About how many paid miles do you expect to run during your next pay week?* We interpret this question as asking drivers for their subjective mean.[11] There are several potential concerns with

---

[8]We use the term "data subset" to distinguish it from the full sample of drivers at Firm A (for whom there is regular personnel data, but no beliefs data) who are studied in Hoffman and Burks (2017) and Burks et al. (2015).

[9]Appendix A.14 describes several unrelated papers using the data subset (e.g., comparing social preferences of truckers, students, and non-trucker adults). Burks et al. (2013) analyze new truckers predicting their quintile on an IQ test to test between different theories of relative overconfidence (people tending to overestimate how well they do compared to other people). Our paper differs from Burks et al. (2013) in that we study absolute overconfidence instead of relative overconfidence; we study beliefs about productivity instead of about performance on an IQ test; and we study beliefs over time instead of at a single point in time. Also, Burks et al. (2013) is focused on testing between different theories of the *causes* of relative overconfidence across people, whereas our paper focuses on the *consequences* of absolute overconfidence for worker behavior and contract design. Although the papers deal with quite different issues, we view the contributions as complementary. Burks et al. (2008) describe the Firm A data collection in detail. 1,069 drivers took part in data collection during training. We restrict our sample to drivers with a code denoting no prior trucking experience or training, giving us 895 workers we are confident are brand-new to trucking.

[10]The question was sent to drivers on Tuesday in 85% of driver-weeks, with the remainder on nearby days (details in Appendix A.7).

[11]Another possible interpretation is that it is asking drivers for the median of their subjective mile distribution for next week. In the data subset, mean and median miles are almost identical (the median of worker miles per week is

using our beliefs question to predict behavior and study overconfidence:

1. Researchers might worry that beliefs are stated to please others, e.g., drivers exaggerate their productivity beliefs to please their boss. However, in our setting, drivers were informed repeatedly that their responses and participation were never to be shared with the company. That is, driver supervisors would never even know whether a worker participated in the survey, let alone what his responses were.

2. No incentives were used to incentivize accurate belief responses. However, as we discuss below in our field experiment with Firm B (see Section 3.2), we find no evidence that beliefs are different when workers are rewarded for accurate beliefs.

3. There is substantial non-response: the average response rate to the weekly beliefs survey is 28%. A 28% response rate may seem low, but is comparable to that in many nongovernmental surveys. For example, in an influential recent study, Card et al. (2012) find a response rate of 20% in a survey of UC Berkeley employees. In Appendix A.1, we re-do our main structural estimation while performing Inverse Probability Weighting (to account for any differential selection on observed characteristics) and show that it has little impact on our estimates. We also estimate a Heckit selection model using the response rate to prior surveys other than the productivity beliefs survey to form an exclusion restriction. We find consistent evidence that non-response bias is limited and is not an important driver of our conclusions. Appendix A.1 provides additional further discussion regarding non-response bias.

One limitation of personnel data is we generally do not see where drivers go when they terminate. Fortunately, we did an "exit survey" by mail for drivers in the data subset. In drivers leaving the firm, the vast majority are not moving to long-haul trucking jobs. Specifically, about 48% of drivers report moving to a non-driving job or unemployment, and 25% went to a local driving job. Only 12% report moving to a long-haul trucking job, and 15% to a regional trucking job. While the response rate on the exit survey was only about 25%, whether someone responds is uncorrelated with most driver characteristics (see Appendix A.5 for more on the exit survey).

1% less than the mean miles per week). Thus, whether workers reported their mean or median expected miles seems unlikely to matter for the reduced-form or structural estimation. See Appendix A.7 for further discussion of belief elicitation methods, as well as the issue of lumpy beliefs/possible rounding.

Summary Statistics. Panel A of Table 1 present sample means on driver characteristics. The median data subset driver is male, white, and 35 years old. Drivers have very low average credit scores. Of the 88% of drivers with credit scores (12% lack a sufficient credit history to have a score), the mean and median credit scores are 586 and 564, respectively, compared to a median of 723 for the US general population at the time of data collection; further, 53% of drivers have a credit score below 600 (i.e., "subprime"), compared to 15% of the US population (Appendix A.4).

Panel B provides quantiles of productivity and productivity beliefs for our main sample, as well as for the sample of 699 drivers used to estimate the structural model. That productivity beliefs exceed productivity on average is easily observed in these simple statistics. In our estimation sample, the median productivity belief corresponds with roughly the 75th percentile in the distribution of actual productivity.

3 Reduced Form Analysis

In this section, we show that, while subjective beliefs are predictive about actual productivity and employee turnover, workers also exhibit a tendency to overpredict productivity. We first present our main results from about two years of non-incentivized belief data from Firm A, and then present the Firm B incentivized data to show the results are robust to incentives.

3.1 Firm A Data

Predicting Productivity. Table 2 shows that beliefs help predict productivity beyond other predictors. We estimate:

$$y_{i,t} = \alpha + \beta b_{i,t-1} + \gamma \bar{y}_{i,t-1} + X_i \delta + \epsilon_{i,t} \tag{1}$$

where $y_{i,t}$ is driver i's productivity in his t^{th} week with the company; $b_{i,t-1}$ is his subjective belief about his productivity in week t stated in week $t-1$; $\bar{y}_{i,t-1}$ is lagged average productivity to date; and X_i are controls. Column 2 estimates $\hat{\beta} = 0.15$, meaning, a driver whose expectation is 100 miles higher than another driver will end up driving an average of 15 miles more. Once average productivity to date or driver fixed effects are added, the coefficient drops to between 0.07 and 0.08. That is, *within person*, subjective beliefs have some predictive power, but less so. Overall, the results suggest that subjective beliefs have informational content, being somewhat predictive across individuals, and

mildly predictive within individuals. The relatively low coefficients likely reflect attenuation bias due to measurement error in subjective beliefs (which we account for in the structural model).

Predicting Quitting. Table 3 shows that quitting decisions reflect subjective beliefs outside of predictors in a standard Bayesian model. We estimate Cox proportional hazard models of quitting of the form:

$$log(h_{i,t}) = \alpha_t + \beta b_{i,t} + \gamma \bar{y}_{i,t} + X_i \delta, \tag{2}$$

where h_{it} is the quit hazard of driver i with t weeks of tenure and α_t is the log baseline hazard. Average productivity to date, $\bar{y}_{i,t}$ is a sufficient statistic for beliefs about productivity in a standard Bayesian normal learning model. However, a 100 mile increase in subjective miles predicts a 6% decrease in the probability a worker quits. The true effects are likely higher, with observed estimates biased downward due to measurement error. The coefficient on beliefs does not change very much as controls are added.

The finding that having higher productivity beliefs is associated with a lower chance of quitting is robust. To show that the finding is not driven by outliers, Appendix Table E1 repeats Table 3 using a dummy for beliefs being above the median (as opposed to a continuous measure) and finds sizeable impacts. The results are broadly consistent when we use lagged beliefs (Appendix Table E2) or a worker's average belief to date (Appendix Table E3), the latter which aims to measure beliefs more of as a stable worker characteristic. These two checks help assuage the concern of reverse causality (e.g., one concern is that people who expect to quit in the future might believe that they will slow down and drive fewer miles). These two checks yield the same result.[12] In contrast to a "standard" setup where workers hold the same beliefs given their productivity signals, workers' heterogeneous subjective beliefs predict quitting.

Overprediction. Although beliefs are predictive, Figure 1 shows that average beliefs consistently exceeds average productivity. Productivity and beliefs are collapsed by week of tenure and then smoothed using a local polynomial regression. Workers initially overpredict by roughly 500

---

[12]Despite the different checks and despite the fact that the coefficient on beliefs does not change much by including observable variables, it is possible that there may be selection on unobserved variables. Thus, we interpret the results here as evidence that overconfidence correlates with fewer quits instead of that overconfidence causes fewer quits. In addition to these checks, one might also think to include some version of (Beliefs minus Productivity) instead of Beliefs as a regressor. However, as we discuss after Prop. 1 in Appendix C, it is the level of a person's perceived inside and outside options that affects quitting in theory, not overprediction. Further, including (Beliefs minus Productivity) as a regressor imposes the restriction that coefficients on Beliefs and Productivity are the same.

miles per week, about 25% of average productivity. This difference declines over time, though it is persistent and decreases very slowly. Even after 100 weeks, worker overprediction is still around 150-200 miles per week. Panel (a) shows means, whereas panel (b) shows medians. Two concerns with Panels (a) and (b) are (i) The sample changes over time (due to quits) and (ii) The productivity line is based on all workers whereas the belief line is based on workers who respond to the survey in a given week. To address (i), we restrict the sample to workers who are there for most of the sample period (at least 75 weeks) in Panel (c). To address (ii), in Panel (d), we re-make the picture dropping the 38% of workers who never respond to the survey. We restrict to workers who are there at least 75 weeks, and look at medians instead of means (to verify results are not driven by outliers). In both cases, the overall pattern of overconfidence is similar, though standard errors are larger.[13]

The average results mask that beliefs are sensible in several ways, and there is a lot of heterogeneity within and across drivers. First, beliefs exhibit aspects of Bayesian updating. As seen in Appendix Table E4, increases in average productivity to date are associated with substantial increases in future beliefs, both across and within drivers. Also consistent with Bayesian updating, in predicting beliefs, the weight on average productivity to date increases with tenure (column 3 of Table E4). Second, although beliefs exceed miles in almost every week when averaged over all drivers, individual beliefs exceed miles only 65% of the time; in 35% of driver-weeks, drivers underpredict, so it's not the case that each driver overpredicts each week. Third, drivers differ substantially in average overprediction. Appendix Figure E3 shows that many drivers are moderately overconfident, some are well-calibrated, and some are very overconfident. Fourth, as mentioned above in Section 2.1, there is a lot of week-to-week variation in productivity that drivers don't control.

That workers' beliefs are sensible in several ways suggests to us that their beliefs are plausible (and not a mark simply of people not taking the survey seriously). Also, Huffman et al. (2018) recently collected data that broadly confirm our results. After our paper first appeared, Huffman et al. (2018) worked with a firm where store managers were asked to predict their quintile in a tournament where store managers competed against one another in terms of quarterly store performance. Huff-

---

[13]In Panels (c) and (d), we stop at 75 weeks instead of the full sample to increase sample size. However, results are similar if we restrict to workers who are there for 100 weeks. More generally, we have made the basic graph comparing productivity and productivity beliefs a number of different ways, including varying means vs. medians, restricting to workers with different tenure levels, restricting based on survey response (all subjects, excluding subjects who never respond, restricting only to weeks where both subject productivity and productivity belief are available), and dropping high outliers in productivity beliefs. Across specifications, although the exact levels of overconfidence vary, the basic graph is broadly similar. In Appendix Figure E2, we plot (Beliefs - Productivity) as a function of tenure, as opposed to plotting beliefs and productivity as separate lines.

man et al. (2018) also find evidence of persistent workplace overconfidence (with overprediction even among managers who had been there for two years or more).[14] Thus, our long-term field evidence on overconfidence in the workplace has recently been replicated in a very different context.

We also emphasize that drivers did not receive direct feedback in the form of someone telling them that they had made overpredictions or underpredictions in the past. Receiving such feedback might have reduced the persistence of overconfidence (Benson and Önkal, 1992).

There are several interpretations of our result that workers tend to systematically over-predict productivity. For example, workers may report aspirations instead of true expectations. Or, workers may report expected miles supposing that "everything goes well" and there are no unexpected hiccups. For both these explanations, one might imagine that misprediction could be eliminated if workers were incentivized to state the mean of their subjective productivity distribution. Alternatively, overprediction may reflect workers' true beliefs and instead reflect a persistent behavioral bias that would be hard to eliminate with an incentive; overprediction may persist, given both substantial idiosyncratic variation in miles and given potential variance bias. To distinguish between these explanations, we turn to incentivized data.

3.2 Incentivized belief data from Firm B

To distinguish between these different explanations and to overcome other concerns with non-incentivized data (e.g., that non-incentivized subjects do not "think hard" enough about their forecasts), we randomized financial incentives for truckers at another large trucking company, Firm B, to accurately guess about their productivity. 272 workers were randomly assigned to guess without financial incentives or to receive up to $10 per week for guessing about their productivity.[15] Subjects did this for about 2-6 weeks before being re-assigned to another treatment: control (nothing

---

[14]Their study is complementary to ours and differs in several important respects. First, while we study absolute overconfidence, Huffman et al. (2018) study relative overconfidence. Second, unlike us (who survey new workers every week for two years), Huffman et al. (2018) survey their managers one time. Third, Huffman et al. (2018) provide evidence in support of selective memory as a mechanism for overconfidence: they ask managers to recall performance in a past tournament (prior to the one being surveyed about), and observe that managers who did well in the past have accurate perceptions, whereas those who did poorly tend to remember themselves doing better than they actually did. This is broadly related to, but different from, our assumptions in the structural model (Section 4 below), where we allow people's perception of signal precision to potentially differ from true signal precision. In our setting, because a majority of people exhibit overprediction, they often get signals that are worse than their beliefs; thus, ignoring or forgetting bad signals is similar to thinking the signal is less precise than it actually is.

[15]While we focus on productivity, we also had subjects guess about their weekly earnings (see Appendix B).

changes), increased incentive (up to $50 per week), or "debiasing."[16] See Appendix B for further information (e.g., how stake size was chosen).

Appendix Table B2 shows that neither the $10 incentive nor the $50 incentive had a significant impact on productivity beliefs. Given that the standard errors are moderately sized, we can't rule out moderate-sized effects in either direction. For example, using column 1 of Table B2, the 95% confidence interval on the impact of the $10 incentive was -131 miles to +66 miles, and the 95% confidence interval on the impact of the $50 incentive was -181 miles to +171 miles. However, given a mean overconfidence level of 250 miles in the column 1 sample, we can reject the hypothesis that *all* the observed overconfidence would disappear if workers were given either $10 or $50 incentives.

Furthermore, Appendix Table E5 shows that there is no evidence that the predictiveness of subjective beliefs toward actual productivity varies with the randomized incentives.

4   Model and Structural Estimation

Our reduced-form analysis suggests that workers overpredict their productivity, as well as that greater beliefs are correlated with a lower chance of quitting. We now develop a structural model of quitting and belief formation to help understand these results, as well as to do a counterfactual of how eliminating overconfidence affects firm profits and worker welfare.

The model is similar to Jovanovic (1979), though it has discrete time and allows for biased beliefs. A worker decides each week whether to quit his job. It is an optimal stopping problem; once he quits, he cannot return.[17] Quitting is the only decision to make—in particular, there is no effort decision.[18] Workers have different underlying productivities, but productivity is initially unknown, both to the worker and the firm. The worker is forward-looking in his quitting decision and each week's miles provides him a noisy signal from which he learns about his underlying productivity. However, workers may be subject to belief biases. The worker's priors need not be accurate, e.g.,

---

[16]"Debiasing" refers to an additional experiment treatment where we provided information about the existence of overconfidence in truckers so as to see if overconfidence could be reduced. Providing information about the existence of overconfidence led to some decreases in productivity beliefs, but impacts seemed to fade with time since treatment, giving us limited power to examine whether randomized changes in beliefs affected quitting. Discussion is left to Appendix B.

[17]While workers who quit Firm A are allowed to re-apply, relatively few return. Of inexperienced workers starting in 2002-2003 who quit during 2002-2003, less than 8% return by the end of 2009.

[18]Effort decisions are not included in most related structural learning models (e.g., Arcidiacono, 2004; Stange, 2012). Our data do not contain exogenous variation in the piece rate that would be needed to plausibly identify the cost of effort function. We speculate, however, that including effort in the model would not qualitatively affect our main conclusions or would actually strengthen them (see Appendix A.8 for further discussion).

he may believe that the job is on average quite lucrative. Further, as new productivity information arrives, he may over- or underweight his prior relative to pure Bayesian updating. In addition to reflecting productivity beliefs, quitting decisions will also reflect a driver's underlying taste for the job or career (e.g., how much a driver dislikes being away from home) as well as idiosyncratic shocks (e.g., a fight with the boss). The firm makes no decisions.[19]

4.1 Model Setup

The time horizon is infinite and given in weeks 1, 2, ... Workers have baseline productivity η, which is distributed $N\left(\eta_0, \sigma_0^2\right)$. Workers are paid by a piece rate, w_t, that depends on their tenure. Workers know the piece rate-tenure profile, and believe that this profile will not be changed by the company at some future date.[20] A worker's weekly miles, y_t, are distributed $N\left(a\left(t\right) + \eta, \sigma_y^2\right)$,[21] and weekly earnings are thus $Y_t = w_t y_t$. $a(t)$ is a known learn-by-doing process, which we specify below. The worker's outside option is r_t and also depends on his tenure. Every period t, the worker makes a decision, d_t, whether to stay ($d_t = 1$) or to quit ($d_t = 0$). Workers make the decision to quit in t having observed their past miles $y_1, y_2, ..., y_{t-1}$, but not their current week miles, y_t. Workers and firms are assumed to be risk-neutral and to have a discount factor given by δ.[22]

Stay-or-Quit Decisions. Workers make their stay-or-quit decisions every period to maximize perceived expected utility:

$$V_t(\mathbf{x}_t) = \max_{d_t, d_{t+1}, ...} E_t\left(\sum_{s=t}^{\infty} \delta^{s-t} u_s\left(d_s, \mathbf{x}_s\right) | d_t, \mathbf{x}_t\right). \tag{3}$$

where $\mathbf{x}_t$ is the vector of state variables ($\mathbf{x}_t$ includes past miles, $y_1, ..., y_{t-1}$, and is detailed further below). (3) can be written as a Bellman Equation: $V_t(\mathbf{x}_t) = \max_{d_t} E_t\left(u_t\left(d_t, \mathbf{x}_t\right) + \delta V_{t+1}(\mathbf{x}_{t+1}) | d_t, \mathbf{x}_t\right)$.

The per-period utility from staying at the job is equal to the sum of the worker's non-pecuniary

---

[19]In the model, the piece rate-tenure profile and training contract are taken as given. In addition, the firm is assumed not to fire workers. In the data subset, quitting is over 3 times more common than firing, and ignoring firing enormously simplifies the model by preventing us from having to estimate a dynamic game.

[20]Assumptions of this form are standard in structural labor and personnel economics, and allows us to avoid having to specify beliefs over possible future firm policy changes. We believe the assumption is reasonable in our setting, given it is not common for the firm to make large changes in the pay schedule.

[21]Assuming that signals are normally distributed is standard in structural learning models (Ching et al., 2013). Visually, the distribution of signals (miles) among all workers has a bell shape centered close to around 2,000 miles, suggesting this assumption is reasonable (and that the distribution is closer to normal than to log-normal or uniform).

[22]Risk neutrality is assumed in many dynamic learning models (e.g., Crawford and Shum, 2005; Nagypal, 2007; Stange, 2012; Goettler and Clay, 2011), though not in all (for examples with risk aversion, see the survey by Ching et al. (2013)). Coscelli and Shum (2004) show that risk parameters are not identified in certain classes of learning models.

taste for the job, earnings, and an idiosyncratic shock:

$$u_t\left(1, \mathbf{x}_t\right) = \alpha + w_t y_t + \varepsilon_t^S,$$

where α is the worker's non-pecuniary taste for the job, and ε_t^S is an i.i.d. idiosyncratic error unobserved to the econometrician (but observed by the worker) with an Extreme Value-Type 1 distribution with zero mean and scale parameter τ. Since workers likely differ unobservedly in taste for the job, we assume there is unobserved heterogeneity in non-pecuniary taste for the job, α, with α drawn from a mass-point distribution (Heckman and Singer, 1984).

If the worker quits, he may have to pay a fine associated with the training contract. Let the vector $\mathbf{k}$ denote the training contract, with k_t the penalty for quitting at tenure t. The utility from quitting is the fine, plus the discounted value of his outside option, plus an idiosyncratic shock:

$$u_t\left(0, \mathbf{x}_t\right) = -k_t + \frac{r_t}{1 - \delta} + \varepsilon_t^Q,$$

where ε_t^Q is an i.i.d. unobserved idiosyncratic error with the same distribution as ε_t^S.[23] Let $V_t^S \equiv E_t\left(u_t\left(1, \mathbf{x}_t\right) + \delta V_{t+1}(\mathbf{x}_{t+1})|1, \mathbf{x}_t\right)$ and $V_t^Q \equiv E_t\left(u_t\left(0, \mathbf{x}_t\right) + \delta V_{t+1}(\mathbf{x}_{t+1})|0, \mathbf{x}_t\right)$ be the choice-specific value functions for staying and quitting, respectively. Plugging in for $u_t\left(1, \mathbf{x}_t\right)$ and $u_t\left(0, \mathbf{x}_t\right)$, the choice-specific value functions are given by:

$$
\begin{aligned}
V_t^Q &= -k_t + \frac{r_t}{1 - \delta} + \varepsilon_t^Q \equiv \overline{V}_t^Q + \varepsilon_t^Q \\
V_t^S &= \alpha + E(w_t y_t | \mathbf{x}_t) + \delta E(V_{t+1}(\mathbf{x}_{t+1})|\mathbf{x}_t) + \varepsilon_t^S \equiv \overline{V}_t^S + \varepsilon_t^S,
\end{aligned}
\tag{4}
$$

and the Bellman Equation can be re-written as $V_t(\mathbf{x}_t) = \max_{d_t \in \{0,1\}}\left(V_t^S\left(\mathbf{x}_t\right), V_t^Q\left(\mathbf{x}_t\right)\right)$.

Agents gradually learn about their productivity as more and more productivity signals are observed. After T periods, we assume that learning about productivity stops. In addition, after this time, there is no more training contract; the worker believes learning is complete; and there is no more growth in underlying productivity, the outside option, or wages.[24] Thus, after passage of T

---

[23]Even though only a portion of the penalties owed were collected, as described in Section 2, we assume that drivers act as if the utility cost of quitting is equivalent to the utility loss from paying the contract penalty. We believe this assumption is reasonable. Firm A was very firm with new drivers about its intention to collect money owed upon a quit. After a quit, drivers who did not pay faced aggressive collection contacts by both Firm A and collection agencies, as well as the reporting of delinquency to credit agencies. As a robustness check, we have experimented with estimating versions of the model assuming drivers act as if the utility loss from quitting is 0.3 times the penalty. Model fit tended to be less good. Indeed, our preferred model still fails to fully match the quitting spike at one year, as seen in Figure 2.

[24]This reflects our focus on weekly outcomes during two years of data as opposed to a very long-term analysis.

periods, worker behavior is governed by the following asymptotic value functions:

$$V^Q = \frac{r_T}{1-\delta} + \varepsilon^Q \equiv \overline{V}^Q + \varepsilon^Q \tag{5}$$

$$V^S = \alpha + w_T E(y_{T+1}|\mathbf{x}_{T+1}) + \delta E(V(\mathbf{x}')|\mathbf{x}) + \varepsilon^S \equiv \overline{V}^S + \varepsilon^S$$

$$V(\mathbf{x}) = \max_{d \in \{0,1\}} \left(V^S(\mathbf{x}), V^Q(\mathbf{x}) \right)$$

Belief Formation. In a standard normal learning model, a worker's beliefs about his period t productivity equals the weighted sum of his prior and his demeaned average productivity to date:

$$E(y_t|y_1, ..., y_{t-1}) = \frac{\sigma_y^2}{(t-1)\sigma_0^2 + \sigma_y^2} \eta_0 + \frac{(t-1)\sigma_0^2}{(t-1)\sigma_0^2 + \sigma_y^2} \frac{\sum_{s=1}^{t-1}(y_s - a(s))}{t-1} + a(t) \tag{6}$$

As t increases, the agent eventually shifts all the weight from his prior to his average productivity signals. We augment the standard learning model in two ways. First, we allow for agents to be overconfident: instead of believing that their productivity, η, is drawn from a distribution $N\left(\eta_0, \sigma_0^2\right)$, agents believe η is drawn from a distribution $N\left(\eta_0 + \eta_b, \sigma_0^2\right)$. Second, we allow for agents to have a perception of signal noise that may be different from the true signal noise: workers perceive the standard deviation of weekly productivity signals to be $\widetilde{\sigma_y}$ instead of σ_y. With these two assumptions, an agent's subjective expectation of his productivity, denoted by E^b (where b stands for belief), is:

$$E^b(y_t|y_1, ..., y_{t-1}) = \frac{\widetilde{\sigma_y}^2}{(t-1)\sigma_0^2 + \widetilde{\sigma_y}^2} (\eta_0 + \eta_b) + \frac{(t-1)\sigma_0^2}{(t-1)\sigma_0^2 + \widetilde{\sigma_y}^2} \frac{\sum_{s=1}^{t-1}(y_s - a(s))}{t-1} + a(t) \tag{7}$$

If η_b is greater (less) than zero, agents exhibit positive (negative) mean bias or overconfidence (underconfidence). As more signals come in, agents learn not to be overconfident, eventually putting zero weight on $(\eta_0 + \eta_b)$. The speed at which this occurs is determined by $\widetilde{\sigma_y}$. Also, since learning is believed complete after T periods have passed, $E^b(y_{T+1}|\mathbf{x}_{T+1}) = \frac{\sum_{s=1}^{T}(y_s - a(s))}{T} + a(T)$.

We allow that workers' reported subjective beliefs include some measurement error, as accurately reporting one's beliefs about productivity may be unusual or unfamiliar for a worker. We assume that reported beliefs equal underlying subjective beliefs plus a normally distributed error. Recall that drivers in week t make predictions about miles in $t+1$. The reported subjective belief, b_{it}, of driver i at tenure week t is distributed: $b_{it} \sim N\left(E^b(y_{it+1}|y_{i1}, ..., y_{it-1}), \sigma_b^2\right)$.[25]

Summary of Within Period Timing. The within period timing in week t is as follows:

1. Workers form beliefs b_t given past miles $y_1, y_2, ..., y_{t-1}$.

2. ε_t^S and ε_t^Q are realized and workers decide whether or not to quit.

---

[25]Note that $E^b(y_{it+1}|y_{i1}, ..., y_{it-1})$ and $E^b(y_{it}|y_{i1}, ..., y_{it-1})$ are the same except for the learning by doing term.

3. y_t is realized, if they do not quit.

Learning by Doing and Skill Accumulation. Productivity increases with the learning by doing function $a(t) = 2a_1 * (-.5 + \Lambda(a_2 * (t - 1)))$, where $\Lambda(x) = \frac{e(x)}{1+e(x)}$ and t is worker tenure in weeks. $a(t)$ depends only on tenure; thus, the speed of learning by doing does not depend on the number of miles driven or on the ability of the driver. Workers fully anticipate the path of $a(t)$.[26]

We also account for skill accumulation following CDL training. After CDL training at Firm A, drivers do "on-the-job training" which includes driving with an experienced driver riding along. We use a length of 5 weeks for on-the-job training.[27] We account for the possibility that drivers may gain valuable skills during this time: we assume the outside option over time is $r_t = r - \frac{6 - min\{t,6\}}{5} s_0$. We fix r using outside data, while s_0, the value of skills from on-the-job training, is estimated. (Besides allowing for skill accumulation during the first 5 weeks, we alternatively estimate the model allowing for continuous skill accumulation: $r_t = r + 2\theta_1 * (-.5 + \Lambda(\theta_2 * (t - 1)))$, where θ_1 and θ_2 are parameters to estimate.)

Solving the Model. The state variables consist of past miles, current tenure, a possible vector of fixed observable individual characteristics (X), a person's taste for the job, a person's level of overconfidence, and the idiosyncratic shocks: $\mathbf{x}_t = (y_1, ..., y_{t-1}, t, X, \alpha, \eta_b, \varepsilon_t)$. The model can allow for heterogeneity in taste for the job and/or in overconfidence. To solve the model, we first solve for the asymptotic value functions (after learning has stopped) using value function iteration. With the asymptotic value functions in hand, backward recursion can then be applied to solve the dynamic programming problem. We provide further details in Appendix D.

4.2 Discussion of Model Assumptions

Outside option. In our model, the outside option, r_t, depends on tenure, but not productivity. This feature differs from many models of firm-sponsored general training where the worker is paid the same share of his marginal product at both his inside and outside option (though less than his

---

[26]The logistic functional form is consistent with Jovanovic and Nyarko's (1996) micro-founded model of learning by doing in which the speed of learning decreases over time, as well as the empirical results on tenure and productivity in Shaw and Lazear (2008). Here, a_1 is the total amount by which productivity increases and a_2 indicates the speed of learning by doing. We believe our assumption that workers fully anticipate the learning by doing process is reasonable in our setting, where the presence of learning by doing seems to be understood.

[27]During this time, drivers often are paid by flat salary instead of by mile. We use a flat salary of $375 per week during on-the-job training. We also assume drivers do not begin learning about their productivity until after 5 weeks, and that non-pecuniary taste for the job is zero during on-the-job training.

full marginal product at both), e.g., Acemoglu and Pischke (1999). We believe our assumption is realistic in our context given that only 12% of workers who exit report moving to a long-haul trucking job, with the vast majority moving to another type of work (see Section 2.2). Having high ability in long-haul trucking does not necessarily imply that one will have high ability in non-trucking jobs or even in short-haul trucking.[28]

In addition, our assumption about the outside option is consistent with an earlier finding that, all else equal, workers with higher productivity to date are substantially less likely to quit. That is, if workers had the same productivity in their inside and outside options and were paid the same share of their marginal product at each, then high and low ability workers would be equally likely to quit. As seen in columns 2-3 of Table 3, this is not the case.

Beliefs. While our model allows for non-standard belief formation, these features are estimated from data instead of imposed. The model does not assume that people have overconfident priors or learn more slowly than would be predicted by Bayes' Rule, but rather these features are identified via the belief data (see Section 4.3); our model nests the standard model as a special case. Several aspects of our generalized normal learning model receive support from the results in Section 3. Differences in subjective beliefs are moderately predictive of differences in productivity across workers, but only mildly predictive within workers. This finding is consistent (broadly) with our modeling assumption that workers do not have private information about their underlying productivity. Underlying beliefs do, however, affect quitting decisions in our model, which is consistent with our earlier empirical finding that, all else equal, workers with greater belief bias are less likely to quit. Although drivers in the model have biased beliefs, they correctly anticipate future changes in their beliefs.

A strong assumption in the model is that workers are not overconfident about their outside option despite potentially being overconfident about their current job ability. However, for over-confidence to "lock in" workers, this assumption is stronger than necessary. Instead, overconfidence will reduce quitting if the worker is more overconfident about his current job earnings than his outside option (see Prop. 1 in Appendix C); i.e., he exhibits *differential overconfidence*. If the strong assumption of no outside overconfidence fails, but workers are still differentially overconfident, over-

---

[28]As discussed in Section 2, in contrast to long-haul drivers who are usually paid per mile, short-haul drivers are usually paid by the hour. For short-haul drivers, relationship-management skills (for managing customer and client relationships) are more important than being able to do a lot of miles quickly while far from home. Some workers will be better at long-haul whereas others are will be better at short-haul.

confidence will still theoretically reduce quitting, but less so than if the strong assumption held. While the assumption of differential overconfidence is difficult to test, we present 6 pieces of evidence and arguments on why it seems reasonable in our setting. Though no piece individually is foolproof, together, the 6 pieces significantly support the assumption of differential overconfidence. We begin with the strongest pieces of argument/evidence, moving down to pieces that are more speculative.

1. Insofar as drivers select the job at which they believe their ability will be the highest, this may lead them toward being differentially overconfident about their ability at the current job relative to the outside option.[29]

2. We collected data on workers' *perceived outside options*. Drivers in the data subset were asked what their earnings would have been had they not started work with Firm A. First, we compare drivers' response to this question to what "similar-looking" people earned in 2006 in the March 2007 Current Population Survey (CPS). As Appendix Figure E4 shows, the perceived outside option workers would have earned had they not gone through training does not appear to be higher than what people like them in the CPS are earning. Second, there is little correlation between a worker's perceived inside option and his perceived outside option (see bolded text in the notes of Figure E4 for details). If workers were equally overconfident about their inside and outside options, we would expect that workers who were more confident about their inside options to also be more confident about their outside options, but this is not the case.

3. Table 3 showed that, all else equal, workers with higher productivity beliefs are less likely to quit while controlling for actual productivity to date. This finding is supportive of differential overconfidence.[30]

---

[29]Van den Steen (2004) provides a "winner's curse" argument on how self-selection can promote belief biases. Consider a worker choosing among several jobs. For each job, he receives a noisy signal about his productivity there. The agent will naturally choose the job with the highest signal, and will be overconfident there relative to other jobs. While it is easy to imagine that workers may have different beliefs when choosing between long-haul trucking jobs and non-long-haul trucking jobs, it is also quite possible that workers might expect to have different productivity at different long-haul trucking firms. For example, one driver may have a lot of experience driving around the South and would be less productive at a firm where most of the routes were in the Northeast. Based on where the routes are, it may be harder for a driver to get home regularly, and spending a lot of time getting home could negatively affect productivity. In a related line of thinking, Lazear (2016) argues that differential overconfidence emerges naturally in a model of occupational choice (since people are likely to choose expectations with positive expectation errors), and provides evidence for this using CPS and PSID data.

[30]In Prop. 1 in Appendix C, we prove that more overconfident workers will be less likely to quit if and only if they are more overconfident about the inside than the outside option. This finding would seem unlikely if workers had the same beliefs about their inside and outside options. Moreover, if having high productivity beliefs was indicative of drivers who think the "grass is always greener" in other jobs, then high beliefs would be correlated with more quitting, not less.

4. While current earnings are proportional to ability in long-haul trucking, most US jobs do not pay piece rates. Thus, even if workers are overconfident about their productivity at their outside option, it is unclear how much this affects workers' perceptions of their earnings at their outside option.[31]

5. It is possible that new workers may form productivity expectations using information from experienced drivers who are reasonably successful as truckers. New workers may fail to fully recognize that those who "make it" as truckers (i) are endogenously selected and that the full sample of new drivers may not be as successful on average and (ii) have already gone through an initial period of increasing productivity.

6. The assumption of differential overconfidence is consistent with evidence in psychology and behavioral economics (see Appendix A.9 for details).[32] The main idea is that, to avoid cognitive dissonance and other psychological discomfort, drivers who have invested significant time, energy, and training contract debt into starting out with working at Firm A may engage in a form of "motivated reasoning" in order to believe that they made the right choice to start work where they did relative to their outside option (Kunda, 1990).

Micro-foundations of Overconfidence. Several theoretical micro-foundations for overconfidence have been proposed in the literature, including evolutionary advantages, self-signaling, and social-signaling. We remain agnostic about the source of the overconfidence (since our model and estimation do not depend on knowing the source). Our contribution is to document overconfidence and explore the implications of overconfidence for behavior and welfare, not to understand its foundation. We do, however, assume that agents do not receive psychological utility from their beliefs, consistent with our Section 3 experimental finding that incentives do not appear to reduce overconfidence. If agents received psychological utility from beliefs, we might expect them to trade off incentives for accuracy with the utility value from stating high beliefs (unless, of course, the personal benefits

---

[31]Performance pay is used in only 37% of U.S. jobs, comprises a median of 4% of total pay across jobs, and is less common in blue-collar jobs like trucking than white-collar jobs (Lemieux et al., 2009). Of course, other pecuniary aspects of a job (e.g., the perceived probability of being promoted to a higher wage) may be affected by overconfidence.

[32]The assumption of differential overconfidence may be more important for the interpretation of the counterfactuals than the structural estimation. In the structural estimation, overconfidence about the inside option varies over time due to learning. Thus, the model-specified overconfidence about the inside option would not be exactly offset by overconfidence about the outside option unless it varied over time in the same way (this is one important way in which the dynamic structural model differs from the one-period model in Appendix C). Appendix A.9 discusses further.

from stating optimistic beliefs are so strong that agents are unwilling to trade-off moderate-sized incentives to reduce their optimism). (To the extent that the assumption is incorrect and truckers do receive substantial psychological benefits from their beliefs, this would cause us to overstate the benefit to workers of eliminating overconfidence.) Also, we remain agnostic whether workers are overconfident about their own skills (e.g., "I have great endurance on the road") vs. whether they are overoptimistic about external events (e.g., "traffic will be better next week").

Learning about productivity. We model quitting as a product of worker learning. This is consistent with Table 3 where, all else equal, workers with higher subjective productivity beliefs are less likely to quit, as are workers with higher average productivity to date. Also consistent with learning, in predicting beliefs, the weight on average productivity to date rises with tenure.

4.3   Estimation and Identification

The model is estimated by maximum likelihood. In Appendix D, we derive the likelihood function and describe the estimation procedure. Although the parameters are jointly identified, we can discuss key data features that allow us to identify particular model parameters.

Productivity and skill parameters. The productivity parameters σ_0, σ_y, and η_0 are identified primarily by the productivity data. σ_0 reflects the degree of permanent productivity differences *across* individuals. σ_y reflects differences *within* individuals in productivity. η_0 reflects the mean average ability of workers in the population. The learning by doing parameters, a_1 and a_2, are identified by how much productivity goes up (a_1) and how quickly (a_2).[33] The skill gain parameter, s_0, is identified based on turnover levels during the first 5 weeks when workers are driving with an experienced driver. The continuous skill gain parameters, θ_1 and θ_2 are identified by how much quitting changes with tenure given the increase in measured productivity.

Taste heterogeneity. The taste for job parameters are identified from persistent differences between individual quitting behavior and the predictions of the model. Suppose that the data contained many low-productivity workers who nevertheless kept choosing not to quit. This would cause the model to estimate that there is a large amount of unobserved taste heterogeneity.

Belief parameters. The subjective beliefs data are critical for identifying the belief parameters. Prior mean bias (overconfidence), η_b, is identified primarily by the difference between believed

---

[33]Note that σ_0, η_0, a_1, and a_2 also appear in E^b. The beliefs and attrition data also help identify these parameters.

and actual productivity, particularly at lower tenure levels. The believed standard deviation of productivity shocks, $\widetilde{\sigma_y}$, determines the subjective speed of learning. The larger $\widetilde{\sigma_y}$ is, the slower that agents' initial overprediction will disappear. The standard deviation of beliefs, σ_b, is identified by noise in beliefs unrelated to information in model-predicted subjective expectations. An increase in σ_b leads to greater week-to-week fluctuations in beliefs unrelated to actual productivity.[34]

Scale parameter. The scale parameter of the idiosyncratic shock, τ, is identified based off of how much quitting behavior in the data differs from that predicted by a model with individual unobserved heterogeneity, but not time-varying uncertainty. Higher levels of τ tend to flatten the quit hazard with respect to worker tenure.

4.4 Implementation

The outside option, r, is taken to be the median full-time 2006 earnings from the 2007 March CPS of workers like the data subset "median" driver (35-year old males with a high school degree), which is \$32,000 per year.[35] We convert this to a weekly wage of \$640. The weekly discount factor, δ, is set to $\delta = 0.9957$, corresponding to an annual discount factor of 0.8.[36] We do not include demographic covariates or heterogeneity in overconfidence.[37] Learning stops after $T = 130$ periods. We use data on up to 110 weeks per driver. After presenting our baseline estimates, we discuss robustness to alternative assumptions. We estimate using 699 workers with complete data (see Appendix D).

4.5 Structural Results

Table 4 displays the main structural estimates and indicates substantial mean bias and variance bias. As a benchmark, column 1 provides estimates assuming no mean bias. Column 2 allows for mean

---

[34]Possible heterogeneity in η_b (discussed in Appendix A.10) is identified from differences across people in the extent of productivity overprediction. We also note that η_b and $\widetilde{\sigma_y}$, in addition to affecting subjective beliefs, will also affect quitting. For example, the faster that agents begin to rely on their average productivity to date in making quit decisions (that is, the faster their quitting decisions reflect "learning"), the smaller that $\widetilde{\sigma_y}$ will be.

[35]In the data subset, "high school graduate" is the modal educational category (40% of drivers) whereas "Some college" is the median category. Table F1 shows our estimates are very similar if we assume a higher outside option.

[36]In most dynamic structural models, the discount factor is assumed rather than estimated, as it is usually weakly identified. An annual discount factor of 0.80 is "low," but is comparable or higher than discount factors used or estimated in other models analyzing dynamic choices of blue-collar or low-income workers (e.g., Paserman, 2008; Fang and Silverman, 2009; Warner and Pleeter, 2001). We have experimented with various discount factors in sensitivity analysis, and assuming a higher annual discount factor such as 0.90 yields quite similar estimates (see Table F1). A discount factor of 0 yields a substantially worse fit, evidence that workers in our context are forward looking.

[37]We have also experimented in the past with model variants that included covariates, e.g., allowing taste for the job, α, to depend on gender, education, race, and age. However, including covariates tended to have little effect on the other parameter estimates and on model fit. See Card and Hyslop (2005) for another dynamic model where covariates are excluded because they do not significantly improve model fit. Appendix A.10 discusses heterogeneity in overconfidence.

bias, estimating bias, η_b, of 674 miles (or roughly one-third of 1,993 miles, the estimated mean of the true productivity distribution). The productivity parameters in column 2 also differ from those in column 1 and seem more reasonable in size.[38] In terms of variance bias, the believed standard deviation of productivity shocks is roughly 2.5 times higher than the actual standard deviation of productivity shocks. This implies that workers update beliefs considerably slower than predicted by Bayes' Rule. Recall that the weight agents place on their signals relative to their prior is $\frac{t\sigma_0^2}{t\sigma_0^2 + \widetilde{\sigma_y}^2}$, where is t is the number of weeks of learning realized following the first 5 weeks of no learning. After 20 weeks of learning, the worker is estimated to place weight 0.38 on his signals (whereas if $\widetilde{\sigma_y} = \sigma_y$, the worker would place weight 0.77 on his signals).[39] Table 4 also indicates significant heterogeneity in non-pecuniary taste for the job.

Table 5 takes the baseline model and adds learning by doing and continuous skill accumulation. The fit is better in both specifications than in Table 4. Many of the parameters are qualitatively similar to before, but there are some differences. In column two, the mean prior bias is larger than before, estimated at 754 miles. The estimated taste heterogeneity is also somewhat different.

Figure 2 shows the model fits the data quite well. We simulate 200,000 drivers using the estimates in column 2 of Table 5. The model tightly matches the survival and productivity-tenure curves. As in the data, the model-predicted quit hazard is inverse U-shaped. The model-predicted quit hazard is initially increasing, reflecting learning about productivity. When workers are uncertain about their productivity, they face an incentive to wait and see how productive they will be before quitting. Also, the model predicts a large spike in quitting after one year (when drivers come off the 12-month contract).

In contrast, model fit without belief bias is considerably worse, as seen in Figure 3. Unlike in Figure 2, the model-predicted survival and productivity-tenure curves do not closely fit the data. Model-predicted average beliefs incorrectly slope up with respect to tenure instead of down as in Figure 2. While the quit hazard has a large spike after one year, one even larger than in Figure 2, the model does far worse predicting the quiz hazard in the first 40 weeks.

Turning from qualitative to quantitative fit, the models with mean bias (column 2 in Tables 4

---

[38] For example, the mean of the prior productivity distribution is 1,993 miles per week, down roughly 20% from 2,436 miles per week in column 1. In column 1, the prior productivity distribution needs to explain the earnings data. But it also needs to explain the quitting and subjective beliefs data, which "pulls" the estimate substantially upward.

[39] This finding is consistent with the well-known psychological phenomenon of "conservatism" (Edwards, 1968), where agents update less than a rational agent would after receiving new information. For recent evidence on conservatism, see Eil and Rao (2011) and Mobius et al. (2014). For theory on conservative updating, see Schwartzstein (2014).

and 5) fit the data much better in terms of overall fit than the models without mean bias (column 1 in Tables 4-5), according to likelihood ratio tests ($p < 0.01$). To analyze model fit on the quitting data alone, we compare the observed weekly number of drivers quitting at week of tenure t, O_t, with the number predicted from the model, E_t, using a χ^2 test. The χ^2 statistic is $\sum_t \frac{(E_t - O_t)^2}{E_t}$. For the baseline models in Table 4, $\chi^2 = 445$ with no bias in column 1, whereas $\chi^2 = 188$ with bias in column 2 ($p < 0.01$). Likewise, for the extended models in Table 5, $\chi^2 = 288$ with no bias in column 1, whereas $\chi^2 = 172$ with bias in column 2 ($p < 0.01$). Thus, the fit in terms of quitting is considerably better in the models with belief bias than without belief bias.[40]

Robustness. Appendix Table F1 shows our baseline estimates are quite robust to different assumptions. Increasing the discount factor does not significantly change the estimates, nor does using inverse probability weighting to address non-response. Winsorizing subjective beliefs at 4,000 miles modestly decreases the mean bias term to 614 miles, which is still quite substantial level of mean bias, suggesting that very high beliefs are not the main driver of the overall mean bias. Allowing learning to occur over 200 weeks (instead of 130) and increasing the outside option also do not much change the estimates. Results are also qualitatively robust to allowing for heterogeneity in overconfidence. Appendix A.10 discusses this and additional robustness checks.

One can imagine alternative non-standard economic forces that affect a worker's taste for the job, e.g., feelings of commitment toward the firm providing training. However, time-invariant shifters of job taste are already accounted for via the taste heterogeneity parameters. In contrast, overconfidence provides a time-varying impact on the value of staying that fits the data well.

Out-of-Sample Fit. Drivers in our sample have the 12-month training contract described in Section 2.1. Another paper (Hoffman and Burks, 2017) studies worker behavior under three different contractual regimes (the 12-month contract, a no contract regime, and an 18-month pro-rated contract), showing that the model developed in the present paper can predict some basic retention patterns under the no contract and 18-month contract regimes. Thus, our structural model also makes reasonable out-of-sample predictions.

---

[40]χ^2 tests are often used to assess the fit of dynamic models (e.g., Keane and Wolpin, 1997; Card and Hyslop, 2005). They can be used to assess the fit of the model with the data, or to compare model fit from two competing models. For example, comparing χ^2 in column 1 of Table 4 with that in column 2 after one parameter is added, the difference in χ^2 is highly significant ($\chi^2_{df=1} = 445 - 188 \rightarrow p < 0.01$). (As caveated in Card and Hyslop (2005), it is more correct to think of the calculated χ^2 statistic as an informal measure of fit, since the predicted numbers are created from the same data being used for the observed cell entries.) To calculate χ^2, we analyze the actual or model-predicted probabilities of quitting in any of weeks 1-110, or in staying through all the first 110 weeks.

5 Counterfactual Simulation: Debiasing

Set-up for Counterfactuals. We use our baseline structural estimates (column 2 of Table 4) to quantify the importance of biased beliefs. Profits are defined as production profits, plus training contract penalties, minus training costs. For a worker who quits in period t, profits are:

$$\pi = \left(\sum_{s=1}^{t} \delta^{s-1}((P - mc - w_s)y_s - FC) \right) + \delta^{t-1}\theta k_t - TC \tag{8}$$

where P is the price the firm charges for one mile of shipment, mc is the non-wage marginal cost per mile (such as truck wear and fuel costs), w_s is the piece rate, y_s is a driver's productivity, FC is fixed costs per week (such as back office support for the driver), θ is the share of the training contract penalty collected by the firm, and TC is training cost per worker. Based on consultation with Firm A managers, we assume that $P - mc = \$0.70/\text{mile}$, $\theta = 0.3$, $FC = \$650/\text{week}$, and $TC = \$2,500$ for the new inexperienced workers we study. We equate the firm's weekly discount factor to the worker's, $\delta = 0.9957$, so as to avoid having results being driven by differences in discount factors; our conclusions are unchanged if we assume higher discount factors for both worker and firm. Further details on computing profits are given in Appendix A.12.1 and Appendix D.

Although workers have biased beliefs, they have standard preferences. Worker welfare is measured by summing earnings, taste for trucking, and idiosyncratic shocks, as in equation (3).[41]

For the counterfactuals, we simulate the full data-generating process for 20,000 simulated workers for up to 1,300 weeks each. While workers are simulated for up to 1,300 weeks, we focus on showing profits per worker and welfare per worker numbers after 110 weeks (corresponding to the maximum number of weeks under observation in the data).[42] We focus separately on profits and worker welfare, and do not analyze total welfare.[43]

Debiasing: Reducing Worker Overconfidence. To examine quantitatively how overconfidence affects quitting, worker welfare, and profits, we simulate eliminating worker overconfidence, which we also refer to, following the psychology literature, as "debiasing." We also consider eliminating overconfidence by one-half, recognizing that debiasing may be incomplete in practice. This

---

[41]Since workers have biased beliefs, average *experienced utility* will differ from ex ante *expected utility*. We measure welfare using experienced utility; this focus is shared by the empirical work of Grubb and Osborne (2015) and the theoretical framework in Mullainathan et al. (2012). Appendix A.12.2 gives more details on computing worker welfare.

[42]The counterfactuals yield the same qualitative conclusions if we analyze profit and worker welfare after 1,300 weeks.

[43]In considering various counterfactuals, while we found our conclusions on profits and worker welfare to be very robust to different assumptions, we found total welfare to depend more closely on particular assumptions made.

practice of analyzing the impacts of counterfactually eliminating a "behavioral" parameter also appears in work such as Handel (2013).

As seen in Table 6, full debiasing increases worker welfare by about 5% since worker quitting decisions become less distorted by overconfidence. Although the workers exhibit significant turnover in the baseline, turnover becomes even higher when workers are debiased. In the un-debaised simulation, worker retention is 41% at the start of week 60, but this falls to 32% under 50% debiasing, and to 24% under 100% debiasing. Without debiasing, workers tend to interpret low-mileage weeks as repeated instances of "bad luck" (i.e., weeks that are low relative to their prior beliefs). After debiasing, workers' quitting decisions are no longer distorted by having a rosy outlook.

In addition, firm profits substantially decline under debiasing. Under full debiasing, profits per worker decline by about $3,000. Due to the increase in quitting, the firm has less time to make profits from a given worker. Our counterfactual allows us to quantify how much overconfidence affects profits and worker welfare in this setting.[44]

6   Conclusion

Using a sample of newly trained truckers, we find that workers tend to persistently overpredict their productivity (on average), thereby providing robust field evidence on worker overconfidence. The difference between average miles and average beliefs eventually declines, but only very slowly. Higher productivity beliefs are correlated with less quitting while controlling for actual productivity to date. To quantify the importance of biased beliefs for profits and welfare, we structurally estimate a quitting model with potentially biased beliefs. We show that overconfidence increases firm profits: if worker overconfidence was eliminated, profits per worker would fall substantially. Further, overconfidence moderately reduces worker welfare by distorting worker decisions.

Our results parallel several studies in behavioral industrial organization indicating that firms may profitably exploit consumers' biases, focusing instead on the behavioral biases of workers. An

---

[44]This counterfactual takes training contracts as fixed. It is conceivable that firms would optimally adjust contracts (wages and training contract penalties) if workers did not exhibit overconfidence. We explored this extended debiasing counterfactual in an earlier version of the paper. Under this counterfactual of debiasing with optimal contractual responses, the overall conclusions from the baseline debiasing counterfactual remain quite similar. The main difference was that debiasing tended to push firms toward decreasing optimal quit penalties. Even if the training contract is fixed, eliminating worker overconfidence might make workers unwilling to accept the job in the first place without an increase in the wage. If eliminating worker overconfidence was also accompanied by an increase in worker wages, this would further increase the worker welfare benefit of debiasing.

important difference in our setting, besides the identity of the parties involved, is the possibility of a baseline market failure of underinvestment (e.g., Acemoglu and Pischke, 1999). In a second-best world with underinvestment in general training, the existence of worker overconfidence makes training more profitable, potentially increasing the quantity of firm training. Paralleling the work in industrial organization, we find evidence that agents (in our case, workers) are harmed by a behavioral bias, but our results also raise the possibility that it may not necessarily be in workers' interests to have that bias eliminated. Unfortunately, because our data are primarily from one firm, we are unable to calculate how overconfidence affects the share of firms willing to train; thus, we are unable to weigh this benefit against the distortion from overconfidence on worker decision-making. Additional research is clearly called for.

While truckers are well-suited for examining overconfidence about productivity, it is important to highlight that we are focusing on one particular job, and the patterns we document may not necessarily hold in other settings. Future worker should examine whether worker overconfidence occurs in other settings. While piece rate compensation is not shared by most other jobs, piece rate compensation is not necessary for overconfidence to make workers less likely to quit. For example, workers may be over-optimistic about some other aspect of the job, e.g., the probability of being promoted. Though we focus on workers in firms, overconfidence may help entice or "lock in" individuals in other labor market situations, e.g., for the decision of enrolling in college or for occupational choice in general.[45]

Worker overconfidence may be important for many aspects of optimal job design and compensation. For example, when firms can choose to pay flat wages or piece rates, paying a piece rate may be appealing if workers are overconfident since overconfident workers perceive they may earn more than they actually will (Larkin and Leider, 2012). Future work should continue to analyze the importance of worker biases for employee behavior, worker welfare, and firm outcomes. Future work should also seek to better understand sources of overconfidence.

---

[45]Stinebrickner and Stinebrickner (2012) show that college students are initially overconfident about their likely performance in college. There has been popular discussion, particularly related to law schools, that students may be overoptimistic about future job prospects when taking on student loans, e.g., David Segal, "Law Students Lose the Grant Game as Schools Win," New York Times, April 2011 and Liz Goodwin, "Law grads sue school, say degree is 'indentured servitude'," Yahoo News, August 2011. In a different application for workers, the impact of overconfidence about job-finding on optimal unemployment insurance, see Spinnewijn (2015). For work on overoptimism and stock options for nonexecutive workers, see, e.g., Oyer and Schaefer (2005).
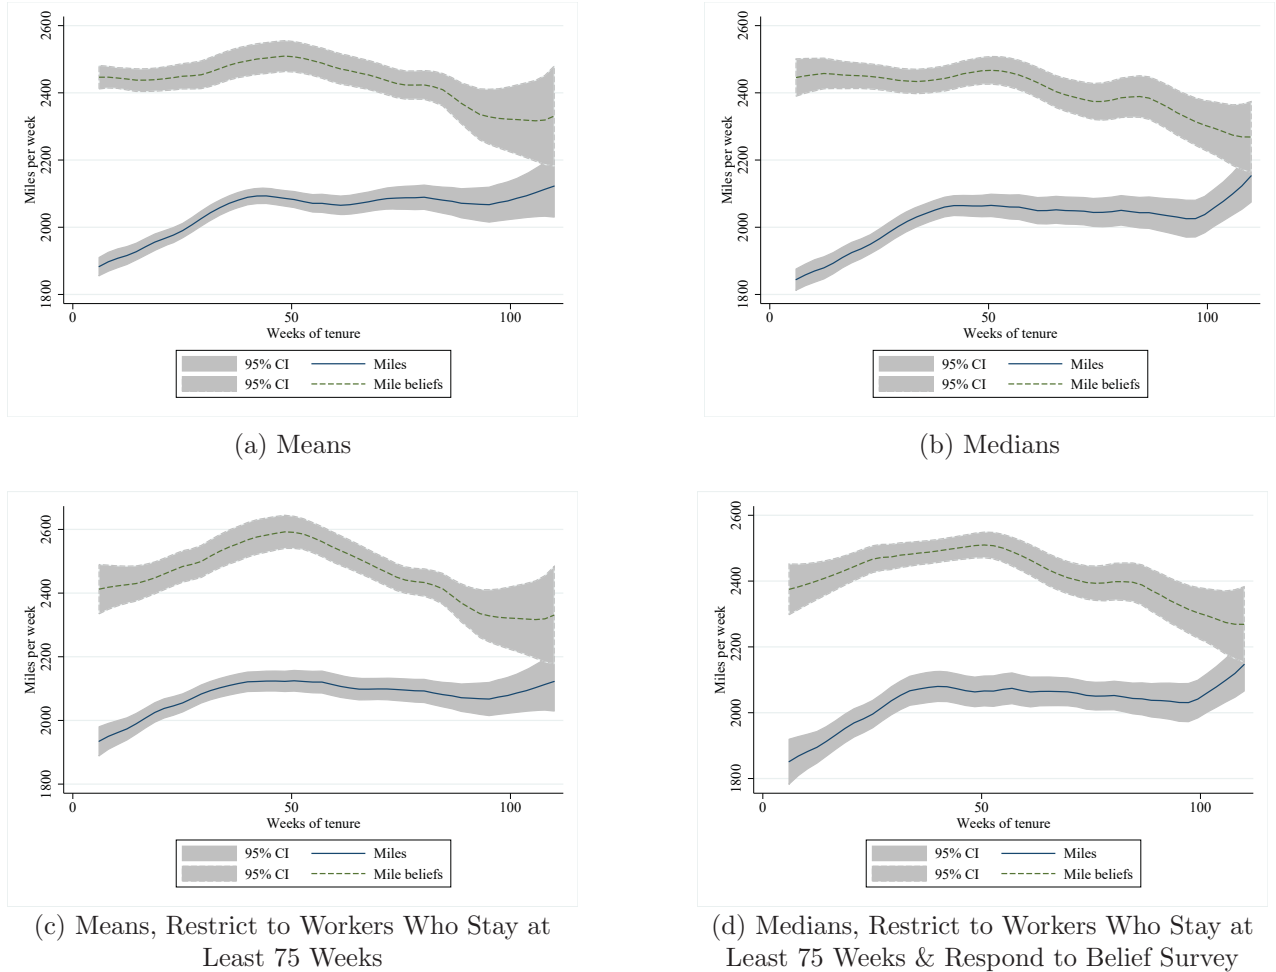
References

Acemoglu, Daron and Jorn-Steffen Pischke, "The Structure of Wages and Investment in General Training," *Journal of Political Economy*, 1999, *107* (3), 539–572.

Arcidiacono, Peter, "Ability Sorting and the Returns to College Major," *Journal of Econometrics*, 2004, *121* (1-2), 343–375.

_ , **V. Joseph Hotz, Arnaud Maurel, and Teresa Romano**, "Recovering Ex Ante Returns and Preferences for Occupations using Subjective Expectations Data," Working Paper 20626, National Bureau of Economic Research October 2014.

Baker, George P. and Thomas N. Hubbard, "Contractibility And Asset Ownership: On-Board Computers and Governance In U. S. Trucking," *QJE*, 2004, *119* (4), 1443–1479.

Barber, Brad M. and Terrance Odean, "Boys Will Be Boys: Gender, Overconfidence, And Common Stock Investment," *Quarterly Journal of Economics*, 2001, *116* (1), 261–292.

Baumann, Andrea, Raisa Deber, and Gail Thompson, "Overconfidence among Physicians and Nurses: The 'Micro-Certainty, Macro-Uncertainty' Phenomenon," *Social Science & Medicine*, May 1991, *32* (2), 167–174.

Bellemare, Charles, Sabine Kroger, and Arthur van Soest, "Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities," *Econometrica*, 2008, *76* (4), 815–839.

Belzer, Michael, *Sweatshops on Wheels: Winners and Losers in Trucking Deregulation*, Oxford University Press, 2000.

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, "Managerial Miscalibration," *Quarterly Journal of Economics*, 2013, *128* (4), 1547–1584.

Benson, P George and Dilek Önkal, "The Effects of Feedback and Training on the Performance of Probability Forecasters," *International Journal of Forecasting*, 1992, *8* (4), 559–573.

BLS, "Bureau of Labor Statistics Occupational Outlook Handbook, 2010-11: Truck Drivers and Driver/Sales Workers," 2010.

Bojilov, Raicho, "Incentives to Work or Incentives to Quit?," 2017. Mimeo, Ecole Polytechnique.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The Value of Hiring through Employee Referrals," *Quarterly Journal of Economics*, 2015, *130* (2), 805–839.

_ , **Jeffrey Carpenter, Lorenz Goette, and Aldo Rustichini**, "Overconfidence and Social Signalling," *Review of Economic Studies*, 2013, *80* (3), 949–983.

_ , _ , _ , **Kristen Monaco, Kay Porter, and Aldo Rustichini**, "Using Behavioral Economic Field Experiments at a Firm: The Context and Design of the Truckers and Turnover Project," in "The Analysis of Firms and Employees: Quantitative and Qualitative Approaches" 2008.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *AER*, 2012, *102* (6), 2981–3003.

_ **and Dean R. Hyslop**, "Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers," *Econometrica*, 2005, *73* (6), 1723–1770.

Chan, Tat, Barton Hamilton, and Christopher Makler, "Using Expectations Data to Infer Managerial Objectives and Choices," 2008. Mimeo, Washington University in St. Louis.

Ching, Andrew T., Tulin Erdem, and Michael P. Keane, "Learning Models: An Assessment of Progress, Challenges, and New Developments," *Marketing Science*, 2013, *32* (6), 913–938.

Coscelli, Andrea and Matthew Shum, "An Empirical Model of Learning and Patient Spillovers

in New Drug Entry," *Journal of Econometrics*, 2004, *122* (2), 213 – 246.

Crawford, Gregory S. and Matthew Shum, "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 2005, *73* (4), 1137–1173.

DeGroot, Morris, *Optimal Statistical Decisions*, New York, NY: Mcgraw-Hill College, 1970.

Edwards, Ward, "Conservatism in Human Information Processing," in B. Kleinmuntz, ed., *Formal Representation of Human Judgement*, 1968, pp. 17–52.

Eil, David and Justin M. Rao, "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *AEJ Micro*, 2011, *3* (2), 114–38.

Fang, Hanming and Dan Silverman, "Time-Inconsistency And Welfare Program Participation: Evidence From The NLSY," *International Economic Review*, 2009, *50* (4), 1043–1077.

Goettler, Ronald and Karen Clay, "Tariff Choice with Consumer Learning and Switching Costs," *Journal of Marketing Research*, 2011, *48* (4), 633–652.

Grubb, Michael D. and Matthew Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," *American Economic Review*, 2015, *105* (1), 234–71.

Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, December 2013, *103* (7), 2643–82.

Heckman, James J. and Burton Singer, "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 1984, *52* (2), 271–320.

Hoffman, Mitchell, "How is Information Valued? Evidence from Framed Field Experiments," *The Economic Journal*, 2016, *126* (595), 1884–1911.

_ **and Stephen V. Burks**, "Training Contracts, Employee Turnover, and the Returns from Firm-sponsored General Training," 2017. NBER Working Paper 23247.

Hubbard, Thomas N., "Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking," *American Economic Review*, 2003, *93* (4), 1328–1353.

Huffman, David, Collin Raymond, and Julia Shvets, "Persistent Overconfidence and Biased Memory: Evidence from Managers," 2018. Mimeo, University of Pittsburgh.

Humphery-Jenner, Mark, Ling Lei Lisic, Vikram Nanda, and Sabatino Dino Silveri, "Executive Overconfidence and Compensation Structure," *JFE*, 2016, *119* (3), 533–558.

Jovanovic, Boyan, "Job Matching and the Theory of Turnover," *JPE*, 1979, *87* (5), 972–90.

_ **and Yaw Nyarko**, "Learning by Doing and the Choice of Technology," *Econometrica*, 1996, *64* (6), 1299–1310.

Keane, Michael P. and Kenneth I. Wolpin, "The Career Decisions of Young Men," *JPE*, 1997, *105* (3), 473–522.

Koszegi, Botond, "Behavioral Contract Theory," *J. Economic Literature*, 2014, *52* (4), 1075–1118.

Kunda, Ziva, "The Case for Motivated Reasoning," *Psychological Bulletin*, 1990, *108* (3), 480.

Larkin, Ian and Stephen Leider, "Incentive Schemes, Sorting and Behavorial Biases of Employees: Experimental Evidence," *American Economic Journal: Microeconomics*, 2012, *4* (2), 184–214.

Lazear, Edward P., "Overconfidence and Occupational Choice," NBER Working Papers 21921, National Bureau of Economic Research, Inc January 2016.

Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent, "Performance Pay and Wage Inequality," *Quarterly Journal of Economics*, 2009, *124* (1), 1–49.

Malmendier, Ulrike and Geoffrey Tate, "CEO Overconfidence and Corporate Investment," *Journal of Finance*, December 2005, *60* (6), 2661–2700.

Massey, Cade, Joseph P Simmons, and David A Armor, "Hope Over Experience: Desirability and the Persistence of Optimism," *Psychological Science*, 2011, *22* (2), 274–281.

Meikle, Nathan L., Elizabeth R. Tenney, and Don A. Moore, "Overconfidence at Work: Does Overconfidence Survive the Checks and Balances of Organizational Life?," *Research in Organizational Behavior*, 2016, *36*, 121–134.

Mobius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya Rosenblat, "Managing Self-Confidence: Theory and Experimental Evidence," 2014. Working paper, UC San Diego.

Moore, Don A. and Paul J. Healy, "The Trouble With Overconfidence," *Psychological Review*, 2008, *115* (2), 502 – 517.

_ , **Samuel A Swift, Angela Minster, Barbara Mellers, Lyle Ungar, Philip Tetlock, Heather HJ Yang, and Elizabeth R Tenney**, "Confidence Calibration in a Multiyear Geopolitical Forecasting Competition," *Management Science*, 2016, *63* (11), 3552–3565.

Mullainathan, Sendhil, Joshua Schwarzstein, and William J. Congdon, "A Reduced-Form Approach to Behavioral Public Finance," *Annual Review of Economics*, 2012, *4*, 511–540.

Nagypal, Eva, "Learning by Doing vs. Learning About Match Quality: Can We Tell Them Apart," *Review of Economic Studies*, 2007, *74* (1), pp. 537–566.

Otto, Clemens A., "CEO Optimism and Incentive Compensation," *JFE*, 2014, *114* (2), 366 – 404.

Oyer, Paul and Scott Schaefer, "Why Do Some Firms Give Stock Options to All Employees?: An Empirical Examination of Alternative Theories," *JFE*, 2005, *76* (1), 99–133.

Paserman, M.Daniele, "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," *Economic Journal*, 2008, *118* (531), 1418–1452.

Sanders, Carl, "Skill Uncertainty, Skill Accumulation, and Occupational Choice," 2016. Mimeo.

Schwartzstein, Joshua, "Selective Attention and Learning," *JEEA*, 2014, *12* (6), 1423–1452.

Shaw, Kathryn and Edward Lazear, "Tenure and Output," *Labour Eco.*, 2008, *15* (4), 704–723.

Spinnewijn, Johannes, "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 2015, *13* (1), 130–167.

Stange, Kevin M., "An Empirical Investigation of the Option Value of College Enrollment," *American Economic Journal: Applied Economics*, 2012, *4* (1), 49–84.

Stinebrickner, Todd and Ralph Stinebrickner, "Learning about Academic Ability and the College Dropout Decision," *Journal of Labor Economics*, 2012, *30* (4), 707 – 748.

Tetlock, Philip, *Expert political judgment: How good is it? How can we know?*, Princeton, 2005.

Van den Steen, Eric, "Rational Overoptimism (and Other Biases)," *AER*, 2004, *94* (4), 1141–1151.

van der Klaauw, Wilbert, "On the Use of Expectations Data in Estimating Structural Dynamic Choice Models," *Journal of Labor Economics*, 2012, *30* (3), 521 – 554.

_ **and Kenneth I. Wolpin**, "Social Security and the Retirement and Savings Behavior of Low-income Households," *Journal of Econometrics*, 2008, *145* (1-2), 21–42.

Wang, James, "Why Hire Loan Officers? Examining Delegated Expertise," 2015. Mimeo, Michigan.

Warner, John T. and Saul Pleeter, "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review*, 2001, *91* (1), 33–53.

Wiswall, Matthew and Basit Zafar, "Determinants of College Major Choice: Identification using an Information Experiment," *Review of Economic Studies*, 2015, *82* (2), 791–824.

Figure 1: Overconfidence: Comparing Subjective Productivity Forecasts with Actual Worker Productivity (as a Function of Worker Tenure)

(a) Means

(b) Medians

(c) Means, Restrict to Workers Who Stay at Least 75 Weeks

(d) Medians, Restrict to Workers Who Stay at Least 75 Weeks & Respond to Belief Survey

Notes: This figure analyzes actual and believed productivity for the drivers in the data subset. Each subfigure is plotted using a local polynomial regression with an Epanechnikov kernel. At a given week of tenure t on the figure, actual productivity is the productivity achieved during that week, whereas believed productivity in t represents the driver's expectation about miles in week $t + 1$. Observations are excluded from the sample if weekly miles are 0, if weekly predicted miles are 0, or if the driver is ever observed in the dataset receiving activity-based pay or salary pay instead of being paid by the mile. We restrict attention to weeks of tenure between 6 and 110 (early weeks involve training and the sample becomes relatively scant after around two years). A bandwidth of 5 weeks is used for the productivity data and a bandwidth of 7 weeks is used for the belief data. In panels (a) and (c), the productivity and belief data are collapsed into weekly means before local polynomial smoothing. In panels (b) and (d), the productivity and belief data are collapsed into weekly medians before local polynomial smoothing. In panel (d), we restrict to the 62% of workers who ever respond to the survey. The figure is similar if we restrict miles to weeks where the driver responds to the survey, though standard errors are larger. It may seem surprising that the initial amount of overconfidence is large. However, as noted in Section 3.1, this figure averages across workers. At an individual level, workers only over-predict their productivity in 65% of driver-weeks (with the percentage calculated excluding instances where miles driven next week is 0, i.e., where the driver is not working). Thus, drivers are not individually over-predicting every week, as there is substantial idiosyncratic variation in miles in the data.

Figure 2: Structural Model: Model Fit
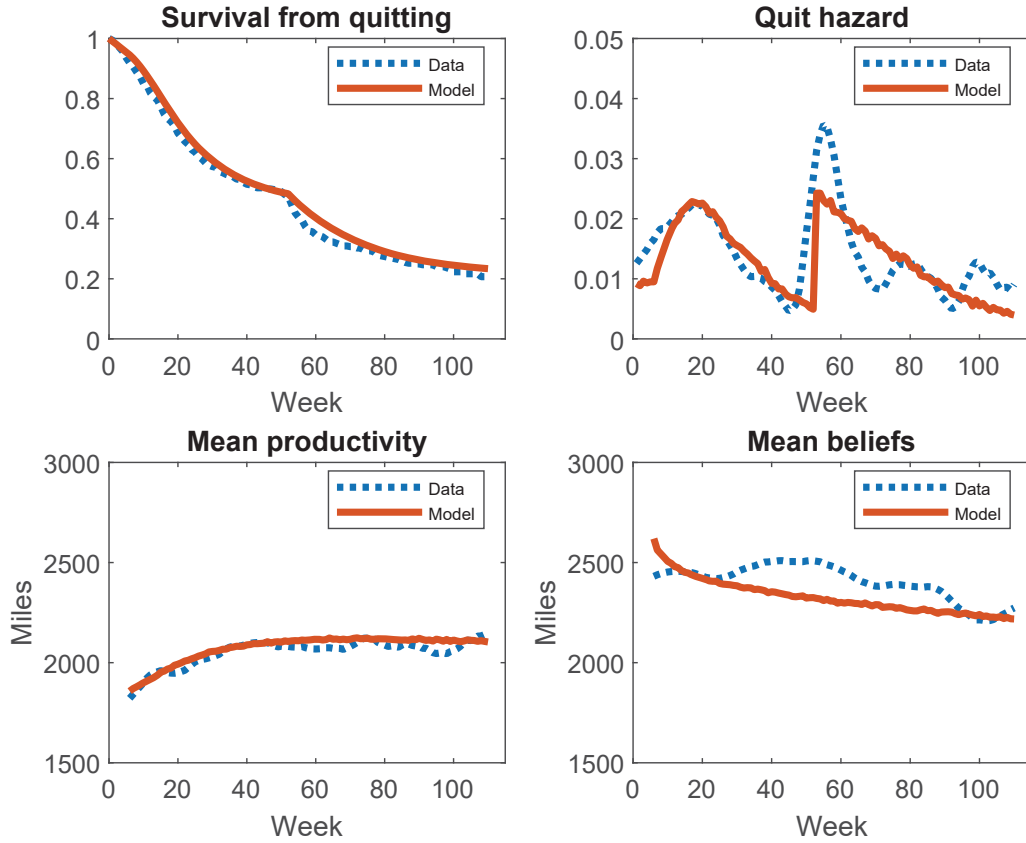
Notes: This figure compares model-simulated data against the actual data to assess model fit. We plot the survival curve, the quit hazard, the mean miles-tenure profile, and the mean beliefs-tenure profile. Survival at week t is the share of workers who survive from quitting to the end of week t. The model simulated corresponds to Column 2 in Table 5. We simulate the entire data-generating process for 200,000 drivers. The data are from 699 drivers with the 12-month contract, all of whom are from the same training school and hired in late 2005 or 2006. The data-based quit hazard, productivity-tenure curve, and beliefs-tenure curve are plotted using an Epanechnikov kernel, with bandwidths of 6 weeks, 5 weeks, and 10 weeks, respectively.

Figure 3: Model Fit, with No Belief Bias
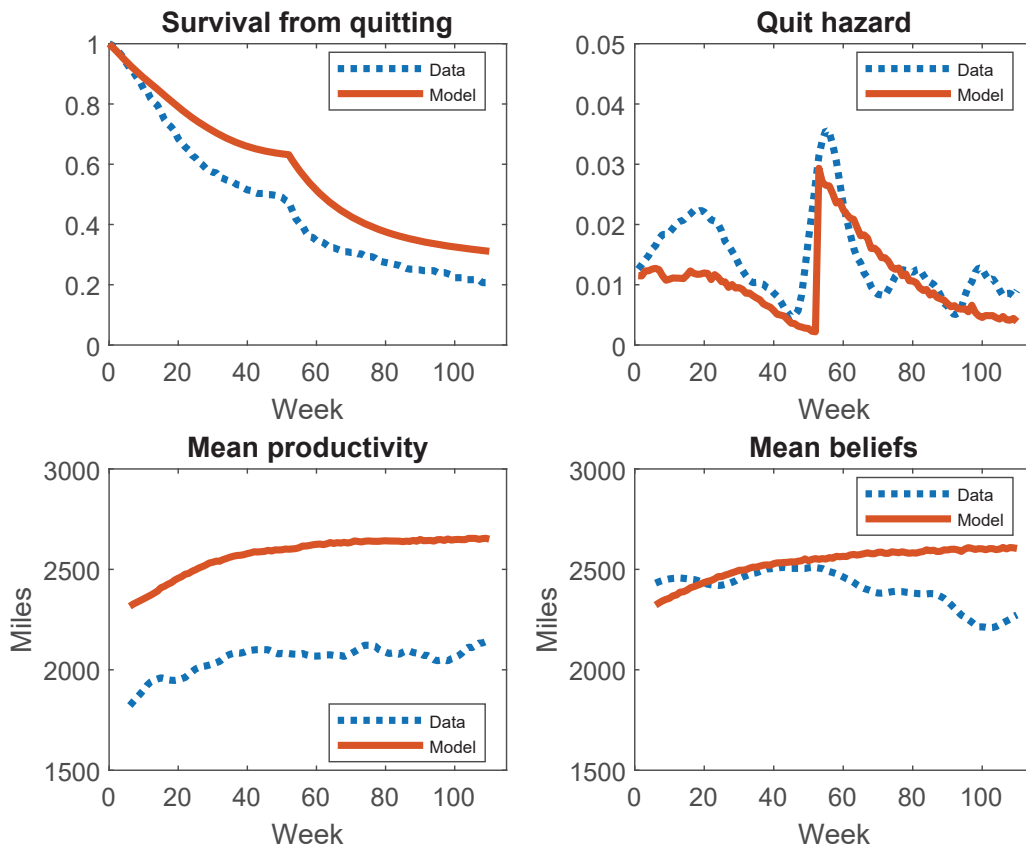
Notes: The notes are the same as for Figure 2 except that the model simulated corresponds to Column 1 in Table 5.

Table 1: Summary Statistics

Panel A: Driver Characteristics

Variable	Mean
Female	0.10
Black	0.11
Hispanic	0.02
Age	36
Married	0.41
Number of Kids	0.96
Years of Schooling	13
Credit Score	586
No Credit Score	0.12
Number of workers	895

Panel B: Productivity and Productivity Beliefs

Percentile	Data Subset		Estimation Sample	
	Miles	Miles Beliefs	Miles	Miles Beliefs
10%	897	1500	1019	1500
25%	1367	1800	1475	2000
50%	1883	2300	1972	2500
75%	2427	2750	2506	2800
90%	2942	3000	3005	3050
Mean	1908	2323	1998	2423

Notes: Panel A provides summary statistics. The drivers in the data are from one of Firm A's training schools and were hired in late 2005 or 2006. Panel B presents quantiles and means on productivity and productivity beliefs, both for the data subset (895 drivers) and for the 699 drivers with complete data that we use in the structural estimation (see Appendix D). Summary statistics on miles are calculated restricting to weeks where miles is greater than zero. See Appendix A.3 for more details on data and sample construction.

Table 2: Do Productivity Beliefs Predict Productivity? OLS Regressions

	(1)	(2)	(3)	(4)	(5)
L. Predicted miles	0.194***	0.147***	0.072***	0.068***	0.080***
	(0.023)	(0.019)	(0.016)	(0.015)	(0.022)
L. Avg miles to date			0.719***	0.623***	
			(0.036)	(0.036)	
Demographic Controls	No	Yes	No	Yes	No
Work Type Controls	No	Yes	No	Yes	No
Subject FE	No	No	No	No	Yes
Observations	8,449	8,449	8,445	8,445	8,449
R-squared	0.070	0.129	0.169	0.191	0.294

Notes: The dependent variable is miles driven per week. An observation is a driver-week. Standard errors clustered by driver in parentheses. Demographic controls are controls for gender, race, marital status, age (dummies for 25-30, 30-35, 35-40, 40-45, 45-50, 50-55, 55-60, 60-80), and education (dummies for high school, some college, or BA plus). Productivity is given in terms of hundreds of miles driven per week. All regressions include week of tenure dummies. The base sample is drivers in the data subset. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 3: Do Productivity Beliefs Predict Quitting? Cox Models

	(1)	(2)	(3)	(4)	(5)
Predicted miles	-0.059***			-0.057***	-0.065***
	(0.018)			(0.020)	(0.021)
Avg miles to date		-0.081***	-0.118***	-0.011	-0.073*
		(0.013)	(0.038)	(0.034)	(0.041)
Demographic Controls	No	Yes	Yes	No	Yes
Work Type Controls	No	Yes	Yes	No	Yes
Observations	8,509	33,374	8,509	8,509	8,509

Notes: An observation is a driver-week. Both predicted miles and average miles to date are in hundreds of miles per week. The regressions are Cox proportional hazard models, where the failure event is quitting. Events where the driver is fired are treated as censored. Standard errors clustered by worker in parentheses. Demographic controls are the same as in Table 2. Column 3 differs from column 2 in that it restricts to the sample to driver-weeks for which there is a corresponding belief expectation. The base sample is drivers in the data subset. In addition, in columns 1, 3, 4 and 5, the sample is restricted to observations with non-missing miles, non-missing average miles to date, and positive mile beliefs. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 4: Baseline Structural Estimates

		No Bias (1)	Belief Bias (2)
	Productivity and Skill Parameters		
η_0	Mean of prior productivity dist	2436	1993
		(18)	(15)
σ_0	Std dev of prior productivity dist	500	292
		(21)	(11)
σ_y	Std dev of productivity shocks	708	708
		(3.5)	(3.5)
s_0	Value of skilled gain wks 1-5	10.6	4.1
		(2.3)	(4.0)
	Taste UH Parameters		
μ_1	Mass point 1 of taste UH	-249	-290
		(8)	(20)
μ_2	Mass point 2 of taste UH	-112	-138
		(15)	(12)
μ_3	Mass point 3 of taste UH	131	145
		(39)	(40)
p_1	Probability type 1	0.55	0.31
		(0.04)	(0.06)
p_2	Probability type 2	0.23	0.46
		(0.03)	(0.06)
	Belief Parameters		
η_b	**Belief bias**		**674**
			(32)
$\widetilde{\sigma_y}$	**Believed std dev of productivity shocks**	**3737**	**1673**
		(163)	(128)
σ_b	Std dev in beliefs	886	877
		(8.2)	(8.0)
	Scale Parameter		
τ	Scale param of idiosyncratic shock	1629	2553
		(180)	(450)
	Log-likelihood	-94401	-94127
	Number of workers	699	699

Notes: This table presents estimates of the structural parameters. The idiosyncratic shock, skill gain, and taste parameters are given in terms of dollars whereas the productivity and belief parameters are given in terms of miles. "Taste UH" stands for unobserved heterogeneity in taste for the job. Standard errors are in parentheses and are calculated by inverting the Hessian. A weekly discount factor of 0.9957 is assumed for workers and firms, corresponding to an annual discount factor of 0.8. The data are from 699 drivers in the data subset, all of whom face the 12-month training contract.

Table 5: Structural Estimates with Learning by Doing and Skill Accumulation

		No Bias (1)	Belief Bias (2)
	Productivity and Skill Parameters		
η_0	Mean of prior productivity dist	2317	1859
		(18)	(22)
σ_0	Std dev of prior productivity dist	498	279
		(22)	(11)
σ_y	Std dev of productivity shocks	708	708
		(1.5)	(1.5)
a_1	Learning by doing level	177	213
		(14)	(14)
a_2	Learning by doing speed	0.10	0.09
		(0.02)	(0.01)
θ_1	Skill accumulation level	101	101
		(31)	(62)
θ_2	Skill accumulation speed	0.02	0.02
		(0.005)	(0.01)
	Taste UH Parameters		
μ_1	Mass point 1 of taste UH	-348	-413
		(36)	(77)
μ_2	Mass point 2 of taste UH	-103	-189
		(31)	(69)
μ_3	Mass point 3 of taste UH	179	204
		(57)	(59)
p_1	Probability type 1	0.64	0.49
		(0.03)	(0.06)
p_2	Probability type 2	0.15	0.28
		(0.03)	(0.06)
	Belief Parameters		
η_b	**Belief bias**		**754**
			(26)
$\widetilde{\sigma_y}$	**Believed std dev of productivity shocks**	**3317**	**1235**
		(158)	**(61)**
σ_b	Std dev in beliefs	883	870
		(1.7)	(1.7)
	Scale Parameter		
τ	Scale param of idiosyncratic shock	1372	2246
		(136)	(282)
	Log-likelihood	-94349	-94043
	Number of workers	699	699

Notes: As in Table 4, the idiosyncratic shock, skill gain, and taste parameters are given in terms of dollars whereas the productivity and belief parameters are given in terms of miles. For the newly added parameters here, a_1 and a_2 are in miles, whereas θ_1 and θ_2 are in dollars. Standard errors are in parentheses and are calculated by BHHH.

Table 6: Counterfactual Simulations

Counterfactual:	Baseline	50% debias	100% debias
Profits per worker	$4,099	$2,747	$1,142
Welfare per worker	$56,856	$58,461	$59,448
Retention at 20wks	0.76	0.59	0.39
Retention at 40wks	0.53	0.40	0.27
Retention at 60wks	0.41	0.32	0.24

Notes: This table reports the results of the counterfactual simulations described in the text, while assuming that training contracts are not adjusted in response. Under the 50% debias and 100% debias counterfactuals, worker overconfidence is reduced by 50% or 100% (by reducing η_b by 50% or 100%). Profits and welfare are defined in Section 5 of the text. The model simulated here corresponds to column 2 of Table 4. "Retention at 20wks" is the share of workers who survive from quitting to the start of week 20.