Making Reference-Dependent Preferences: Evidence from Door-to-Door Sales

Samuel Dodini*
Federal Reserve Bank of Dallas

October 2024

Abstract

This paper uses data from a door-to-door sales company and an online experiment to examine the relationship between reference-dependent daily labor supply and long-run goal achievement. In the sales data, I show that daily labor supply kinks downward at a worker's expectations and that these expectations directly correspond to bonuses paid at the end of the sales season. The bonuses induce workers to adopt long-run targets and subsequently distribute these into internalized daily goals around which they exhibit loss aversion. These dynamics explain why non-linear payment schemes increase performance: workers change their short-run behavior in response to long-run performance targets. The online experiment confirms a causal interpretation of this relationship between bonuses and short-run behavior and supports the idea that short-run reference dependence can be "made" or induced by firms by adopting non-linear compensation schemes. These dynamics increase worker output and firm profitability and can explain why non-linear compensation is so popular in the labor market.

Keywords: reference dependence, loss aversion, non-linear compensation, goals
JEL Codes: D9, J22, J33, M52

*Email: samuel.dodini@dal.frb.org. I am grateful to Michael Lovenheim, Maria Fitzpatrick, Evan Riehl, Ted O'Donaghue, Alex Rees-Jones, Seth Sanders, Sam Hirshman, Bertil Tungodden, Linh Tô, and Devin Pope for helpful comments, guidance, and support with the experimental design. Thank you to Sebastian Fest for support in implementing the online experiment. I am also thankful to seminar participants at the Cornell Behavioral Economics Workshop, Cornell Labor Economics Workshop, and Cornell Labor Work in Progress Seminar for helpful comments and discussion. This project was partially funded by the Research Council of Norway through its Centers of Excellence Scheme, FAIR project no. 262675. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 Introduction

People often fail to reach their personal performance targets due to motivational problems (Dellavigna, 2009). From a firm's perspective, this is particularly problematic for the firm's objectives. One possible approach to addressing these issues is to subdivide a target into components and evaluate one's performance in a smaller window or "narrow bracket" such as a single task or a shorter time frame. In other words, people may set a short-run goal. These goals then act as reference points (Heath et al., 1999).

The motivational power of narrow goal setting relies upon people behaving in ways consistent with reference-dependent preferences (Heath et al., 1999; Koch and Nafziger, 2016; 2020; Imas et al., 2017). That is, a person's utility will depend not only on her absolute performance but also on where her performance stands relative to some mental target, i.e., the goal. Under loss aversion, when the person is operating below that mental target, the negative comparison to her target induces lower utility. However, her marginal utility is higher in this state. All else equal, this leads her to exert more effort until she reaches the target. Thus, by creating reference points, sophisticated behavioral agents can use self-imposed, psychological costs in the short term to overcome self-control problems in the long term. From a firm's perspective, it is useful to induce workers to such behavior if they suspect self-control problems and dynamic inconsistency might arise.

The idea of using short-term goals as reference points in the service of larger objectives is intuitive. However, there is limited field evidence of what this looks like in real-world labor market settings where workers interact with firms and where one would expect to find such dynamic inconsistency most abundant. It is also unclear why workers might choose to adopt short-run reference points in the absence of direct daily incentives, as they appear to do in the empirical literature on taxi drivers that has come to define much of our real-world understanding of short-run reference dependence in labor markets (e.g. Camerer et al. (1997); Crawford and Meng (2011); Thakral and Tô (2021)). Rather than being a cognitive bias or error, do workers leverage daily goals as a commitment device to achieve longer-run targets? And importantly, can firms induce workers to take up this strategy at a relatively low cost? Despite the fact that worker-firm interactions represent the most consistent ways in which performance targets are adopted in a person's life and where people are most likely to shirk, we know very little about how such short-run references are "made", shaped, or induced by firms, nor is there clarity about how this might relate to the firm's and the worker's longer-run objectives.

This paper investigates reference dependence and goal setting in two contexts: a context new to the literature, door-to-door sales, and an online real-effort task experiment. The main contribution of this paper is to establish two important, new empirical observations in both contexts: 1) workers endogenously exhibit reference dependence in their short-run labor supply and effort choices as a means of holding themselves accountable to their longer-run objectives even in the absence of a loss-framed contract; and 2) a firm's choice of compensation scheme (i.e. non-linear compensation via discrete performance bonuses) can causally induce workers to "make" such short-run references by making particular longer-run targets salient. This action cost-effectively increases worker performance and firm profits. This is consistent with theories of goal-setting in which workers use loss aversion as a commitment device in the short run, even when contracts are not loss-framed and even in the absence of short-run monetary incentives. Bonuses, therefore, induce rational short-run reference dependence to address expected present bias—using one behavioral bias to overcome another. This effect has never been demonstrated before in the literature. This result at least partially explains why non-linear compensation schemes such as bonuses are so popular in the labor market.

I first motivate my analysis with a discussion of the theories behind reference dependence and goal-setting behavior and the mechanisms explaining why workers might prefer loss-framed contracts or impose short-run loss aversion upon themselves. I draw upon the recent Koch and Nafziger (2020) model, which builds on the fact that present bias leads to suboptimal short-term effort because present-biased agents will tend to shirk today in favor of expecting to work harder tomorrow. In short, the model explains how reference-dependent preferences and loss aversion in the short run might solve self-control problems, even when a long-run target is not a loss-framed contract (e.g. Imas et al. (2017)) and even without short-run monetary incentives (e.g. Kaur et al. (2015).[1] I propose that firms can influence workers to adopt loss aversion around internal, short-run goals by 1) making long-run targets more salient and 2) attaching significant monetary consequences to long-run output. These raise the incentive for workers to strategically adopt narrow brackets and short-term loss aversion in order to achieve their more visible and salient long-run output if they believe their future self may be tempted to shirk. I use my sales data and experiment to examine these key dynamics.

I analyze high-frequency data from a sales company that employs fixed-term, commission-based contractors. First, I establish the baseline observation that workers do exhibit reference-dependent labor supply in this novel setting. I test for this on two margins: the extensive margin (the choice to stop working for the day) as well as the intensive or "exertion" margin (effort conditional on working). I provide clear evidence that door-to-door sales workers exhibit loss aversion around expectations in their extensive margin labor supply choices.

I use a detailed panel of observations in half-hour increments with each seller's location, cumulative service contracts generated (which I call "sales" throughout the paper), pitches presented to a prospective customer, and the probability of stopping work for the day (the extensive labor supply margin). My measure of intensive margin effort (or alternatively phrased "exertion") is pitches per half hour. I define the reference point as a sample proxy of expectations: each seller's

[1]The model stands in contrast to the standard model, which assumes rationality in intertemporal utility affecting long-run performance (i.e., no self-control problems and, therefore, no need for short-run goals) and no utility responses to reference points in the short run. Under the behavioral model with a non-zero degree of present bias, firms find it preferable to encourage narrow brackets.

own average daily number of sales for all past workdays in the season, which I show is highly correlated with revealed long-run objectives. Upon reaching their expectations-based reference point, the probability a worker stops for the day increases significantly by a factor of 2.8–4.1 times relative to below the reference point, suggesting that losses loom larger than gains by a factor of approximately 3 to 4. On the "exertion" margin, the change is quantitatively small. The choice of when to stop working is the key margin at which reference-dependent daily labor supply operates.

I then examine the relationship between a worker's daily sales, work hours, and the firm's lump-sum bonuses paid at the end of the season. The commitment device hypothesis suggests that the firm's contract structure incentivizes the worker to optimize around a long-run goal at a bonus threshold and workers then distribute that goal into daily targets. I show three pieces of evidence to support this hypothesis. First, I show that these workers are forward-looking, as evidenced by the fact that their work hours do not significantly respond to the increases in commission rates that follow their cumulative sales. Rather, they set plans for their work hours based on their anticipated commission rates. Second, I show that the distributions of performance are subject to significant "bunching" around bonus thresholds, and this emerges early in the sales season. This is most evident in the distributions of performance among those targeting the same bonus thresholds.

Third, I show that upon reaching their relevant bonus threshold, workers significantly reduce their work hours despite continuing to be paid a slightly higher piece rate than they were paid before reaching the bonus. Thus, attainment of the bonus appears to be the motivating factor behind work hours persistence and daily reference dependence and not other mechanisms like simple habituation. These three observations together provide strong evidence that the firm, through its compensation scheme, can shape the choice of long-run objectives and induce short-run targeting.

To bolster the interpretation that the relationship between short-run reference dependence and the bonus schedule is causal, I conduct an online experiment designed to mirror in a short period the dynamics of the sales data. In the experiment, participants in a multi-round task are randomly offered either a piece rate or a bonus upon reaching a certain performance level, where the bonus is only evaluated at the end of the fourth/last round. I then observe the distribution of performance in each round. Significant excess "bunching" of the performance distribution each round in the bonus treatment provides clear evidence that even though round-specific performance is not payoff-relevant for those in the bonus payment condition, participants exhibit reference dependence with loss aversion around endogenously self-selected targets. In other words, they anchor their performance to specific targets each period to ensure that they satisfy the bonus condition at the end of the task. Many respondents make this strategy explicit in open-ended responses. Upon reaching the bonus threshold in the final round, participants in the bonus condition significantly reduced their effort. Despite this, participants in the bonus treatment significantly outperformed those in the piece rate treatment: for the same average performance across both conditions, the bonus group incurred a 31% lower per-person compensation cost. These patterns follow closely the patterns in the sales data.

The core contribution of this paper is to bring together real-world and experimental evidence that short-run reference dependence acts as a commitment device to achieve long-run goals and that firms can induce such behavior to improve worker performance. In doing so, the analysis contributes to three main strands of the literature on reference dependence.

In the first strain, several papers analyzing taxi and rideshare drivers' behaviors have found a negative relationship between daily wages and hours worked, downward shifts in labor supply at particular earnings levels, or negative labor supply responses to large tips, consistent with reference dependence.[2] However, the taxi cab literature has not empirically explored the purpose of having reference points at all—adaptive or fixed. Even though the earliest taxi cab studies (e.g. Camerer et al. (1997)) hypothesized that income targeting may help drivers address self-control problems, none of these studies have empirically explored this dynamic. In a recent experiment in Kenya, Dupas et al. (2020) show that a person's stated income needs and expectations for earnings (rather than just total income) act as reference points. The authors suggest such targeting motivates workers as a commitment device to perform their physically demanding jobs, though intertemporal dynamics are not explored as they do not consider measures of broader income needs.[3] This empirical literature has not considered the interaction between clearly defined long-run objectives and daily targets in reference dependence. This paper makes a significant contribution by examining both the short- and long-run in a unified way using real-world data and a controlled experiment.

The "lumpy" nature of income in my sales context and the lottery-like nature of success at each door that decouples immediate income from effort makes this setting quasi-experimental and ideal for the study of loss aversion.[4] In addition, previous studies testing reference dependence are quite narrowly focused on routine and manual tasks like taxi driving or physical labor. This paper provides new evidence in both a simple manual setting (the online experiment) and a novel work context that uses adaptive cognitive and social skills in a developed country. Understanding this skill distinction is crucial if workers in manual occupations differ significantly in their attributes from those who select into primarily social or cognitive occupations or who have the education to

[2]See Camerer et al. (1997); Chou (2002); Crawford and Meng (2011); Farber (2015); Morgul and Ozbay (2015); Agarwal et al. (2015); Martin (2017); He et al. (2018); Schmidt (2018). Other papers that find evidence of reference-dependent labor supply analyze the behavior of bike messengers (Fehr and Goette, 2007; Goette et al., 2004) and fishermen in Hawaii (Nguyen and Leung, 2013). However, the literature is far from settled. A competing set of studies of drivers finds a positive relationship between daily wages and hours worked and concludes that the standard model performs better than prospect theory (Farber, 2005; 2008; 2015; Sheldon, 2016). Other analyses that find evidence supporting the standard model examine day laborers in Malawi (Goldberg, 2016), stadium vendors (Oettinger, 1999), fishermen in Florida and India (Stafford, 2015; Giné et al., 2016), and markets in India (Andersen et al., 2014).

[3]Their questionnaire and definition of "income needs" is specific to each day's idiosyncratic needs, which often exceed their average income.

[4]Encountering one extra resident willing to purchase the sellers' services leads to an increase in income of \$100–\$250. An extra sale or two by a seller is worth the same amount as an entire shift for a taxi driver (\$270) but takes roughly the same amount of time as 1-2 taxi trips (16–32 minutes) (Thakral and Tô, 2021).

enter these occupations.[5]

The second strain focuses on reference dependence as expressed in the distributions of final outcomes around a single or ending target. In a firm-worker setting, Kuhn and Yu (2021) examine the effects of kinks in a commission schedule on final team performance and find these act as symbolic rewards, leading to bunching in the distribution of performance. Cai et al. (2022) examine kinks in the compensation schedule at a Chinese manufacturing firm to estimate labor supply elasticities and firm cost savings. Their analysis does not consider behavioral factors such as loss aversion. Similarly, Freeman et al. (2019) analyze a shift in a performance bonus threshold at a Chinese insurance company, which led to significant increases in worker output and firm profits. Beyond knowing that these incentive schemes increased total worker output at the end of an evaluation period, little is known about how or why these approaches were effective.[6] The underlying day-to-day behavioral dynamics have immense implications. If a new compensation scheme induced workers to set short-run goals with loss aversion, this represents, from the firm's perspective, a low-cost *psychological* incentive rather than a high-cost *monetary* incentive. This analysis documents this effect in two distinct settings.

The third strand experimentally ties the adoption of commitment devices that leverage loss aversion to increase a worker's output. For example, laboratory and field experiments show that those instructed to set daily goals perform better than those instructed to set weekly goals (Koch and Nafziger, 2020), that workers prefer loss-framed contracts (Imas et al., 2017), and that many workers will voluntarily reduce their daily compensation if they fail to meet daily targets when offered these dominated contracts (Kaur et al., 2015). My two settings provide distinct contributions relative to this work. First, short-run goals in my setting are not exogenously assigned but are *endogenously* adopted by workers in response to the structure of longer-run incentives. Second, it is through this endogenous adoption of narrow brackets that performance increases, thus providing a dynamic mechanism for the performance increases found in Imas et al. (2017), Kuhn and Yu (2021), and Freeman et al. (2019). I show that such a mechanism can be induced by firms at a low cost, even without a loss-framed contract and without the need for strong daily performance incentives (as in Kaur et al. (2015)), as these may be costly to monitor and implement for the firm.[7] Thus, this paper contributes significantly to our understanding of the interaction between

[5]Given the literature suggesting significant differences in key behavioral parameters across occupation, education, or cognitive ability (e.g. discount rates, risk preferences, etc.), it is reasonable to suspect possible differences in loss aversion and reference dependence (e.g. Cadena and Keys (2015); Bellemare and Shearer (2010); Patnaik et al. (2020); Fouarge et al. (2014); Warner and Pleeter (2001)).

[6]Explicit long-run goals act as references, as seen in the distributions of finishing times among marathon runners (Allen et al., 2017; Markle et al., 2018). However, like cumulative outcomes in a firm, marathon times are the result of dynamic processes (e.g. daily training). Despite this, previous studies have not focused on how the long-run target interacts with short-run choices or preferences.

[7]My analysis suggests that if the workers in Kaur et al. (2015) were paid longer-run bonuses for cumulative performance, they might have endogenously chosen daily references and goals and imposed a utility cost upon themselves without needing to be offered the dominated daily contracts. Thus, I also contribute to the literature on "insider

firm-provided incentives and dynamic worker choices and helps explain why non-linear payments are so popular across many industries and occupations.

My analysis of long-run objectives and short-run expectations is the first field study of which I am aware to examine how endogenous period-specific/daily reference dependence acts as a commitment device to achieve revealed or imposed long-run goals in the absence of loss-framed contracts—to address the "why" of short-run reference points. I show that firms can leverage salient long-run targets to induce this behavior and that workers are responsive to firms' broad incentives even in their daily or period-specific activities. Self-imposed, short-run reference points appear malleable when long-run incentives change. My online experiment confirms the causal interpretation of these dynamics by linking an increase in reference-dependent behavior to non-linear compensation schemes. Firms can take advantage of these dynamics because reference dependence makes it easier to motivate a worker if she perceives herself to be in a "loss" domain. The question of reference-dependent labor supply is, therefore, central to our understanding of the power of incentives to motivate effort.

2 Conceptual Framework

2.1 Reference Dependence vs Standard Model for Labor Supply

One key insight of Prospect Theory is that losses loom larger than similarly sized gains (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). The implication for labor supply under loss aversion is that workers supply greater effort while in a loss domain (before achieving a target) relative to what they supply in a gain domain (after achieving a target). This leads to a discontinuous change in marginal utility after surpassing some reference point, with the marginal utility of income falling significantly by some factor $1/\lambda$, where λ is the parameter of loss aversion. This induces a discontinuous change in labor supply, all else equal. Importantly, no such discontinuity is predicted by the standard model. Because loss aversion is costly due to the agent experiencing lower levels of utility while in the loss domain, this behavior is sometimes construed as a negative cognitive bias. However, under certain conditions, loss aversion may be rational.

The location of the reference point that determines the gain and loss domains is important. In an essential theoretical paper (Kőszegi and Rabin, 2006), the KR model theorizes that "recent expectations" act as important reference points. The KR model proposes that these expectations are determined in what they call "personal equilibrium," that is, by behaviors that are optimal given the worker's expectations about the future. Put another way, a forward-looking worker can make a plan around what she perceives to be the optimal path forward, and when the final choice is made in real-time, the planned path becomes her reference point. Kőszegi and Rabin (2006) posit when introducing this theory that firms can play a significant role in establishing a worker's expectations;

econometrics," the use of non-linear incentives and bonuses, and their effects (Ichniowski and Shaw, 2003).

however, empirical evidence for this role is generally sparse.[8]

2.2 Loss Aversion, Goals, and Short-Run References

Multiple theoretical treatments propose that the strategic use of loss aversion may be rational if a worker has a problem with self-control or dynamic inconsistency as a result of present-biased preferences (see Ariely and Wertenbroch (2002); Shefrin and Thaler (1992); Camerer et al. (1997); Imas et al. (2017); Kaur et al. (2015), with Koch and Nafziger (2016; 2020) presenting the most formal recent treatment). Present bias has been documented in a variety of contexts, e.g. exercise goals (DellaVigna and Malmendier, 2006), education (Ariely and Wertenbroch, 2002), credit markets (Meier and Sprenger, 2010), and savings (Ashraf et al., 2006), and it is simple to extend the concept to labor markets. Put briefly, time-inconsistent preferences lead to effort levels in the current period that fall well below what a worker would have chosen for herself ex ante.[9]

It is worth explicitly exploring this dynamic under the framework in Koch and Nafziger (2020). Suppose workers perform the same task each period or day (in time $t \in [1, T]$) with effort level e_t incurring costs $c(e)$ that are convex. Then suppose there is a total benefit b at the end of a long-run evaluation period that is a function of total effort, and effort is deterministic over utility outcomes. If a worker is a quasi-hyperbolic discounter (Laibson, 1997), then there are t versions of the worker, one for each day, with utility $U_t = u_t + \beta[\sum_{\tau=t+1}^{T+1} u_\tau]$ and instantaneous utility u_t and a present-bias factor of β. Instantaneous utility is $u_t = -c(e_t)$, and final period utility $u_{T+1} = \sum_{t=1}^{T} b(e_t)$. Ex ante, a period 0 self sets marginal costs and benefits equal such that $\beta = 1$ and $b'(e_0^*) = c'(e_0^*)$. This would be the equilibrium effort under the standard model of some chosen long-run outcome.

Now suppose each period's self after period 0 discounts future benefits by $\beta < 1$. Equilibrium effort with present-biased preferences would be $\beta b'(e_0^*) = c'(e_0^*)$. A worker who set out to perform at e_0^* to achieve total benefit $b(e_0^*)$ in time 0 has an incentive in time t to substitute effort from today to tomorrow or from the current period to the next. Because total outcomes are fungible across days, substituting effort across days may lead to suboptimal effort in time t under the ex-ante assumption that the worker may increase effort in $t + 1$. Importantly, if the benefit at the end of the period (for example, a total payout for a worker) were increased by some proportion γ, while the worker has an increased incentive to gain benefit $\gamma b(e_0^*)$ at the end, the utility benefit each period would only increase by $\beta \times \gamma b'$. Practically speaking, that means present bias blunts the incentive effect of additional benefits to perform in the longer run or increasing piece rates, making such incentives less cost-effective.[10]

[8] Appendix E provides more general background on linear gain-loss utility in labor supply, as well as expectations and goal-setting.

[9] For a comprehensive discussion of present bias, see Dellavigna (2009).

[10] In in many occupations like sales, effort costs to achieve the same objective may fluctuate through a day-specific cost function ($c_t(e_t)$), that is, the time and effort cost of achieving the same objective. The standard model predicts a

But suppose self 0—a forward-looking agent—sets a narrow bracket through a daily or period-specific goal to bind the incentives for self t in the future through additional comparison utility penalties, i.e. for $e_t < g_t$, $\hat{\beta}(g_t - e_t)$. For a sophisticated individual who correctly predicts β and calibrates $\hat{\beta}$, personal equilibrium suggests that g_t should be the same as the optimal effort that period 0 self would choose given their beliefs about future effort, or in other words, that $\hat{e}_{t,0} = g_t$. When tasks are repeated daily, $g_t = [b(e_0^*)]/T$. Self t then provides effort g_t each period, thus solving the self-control problem.

Combating suboptimal effort substitution is the key incentive introduced by narrow bracketing. In the case of daily or period-specific goal-setting, because the marginal utility of income is higher in the loss domain, workers have the incentive to exert more effort on "high-cost" days to achieve a minimum performance. On "low-cost" days, they surpass their target more easily, but the marginal utility of additional income falls, so the worker has the incentive to reduce their labor supply upon surpassing it. Thus, for the worker, there is a cost to narrow bracketing: negative comparisons in the loss domain reduce experienced utility while in that domain. Therefore, workers using narrow brackets as commitment devices will do so only until the broader goal is reached, after which there is no reason to continue engaging in negative comparisons. Because of these costs and despite the possible positive effects on goal attainment, not all workers may engage in this behavior.

The idea of using short-run mental targets is mentioned as a possible explanation for observed daily income targeting in the first empirical analyses of taxi driver behavior (e.g. Camerer et al. (1997)), though the analysis does not explore it in detail.[11] Dupas et al. (2020) similarly invoke this explanation, though the analysis considers stated daily income needs rather than goals or broad objectives. This is important because "income needs" can apply to any context in which budgetary needs (e.g. rent/mortgage payments) may act as medium- or long-run targets. In the context of, for example, taxi drivers, setting a daily target and exhibiting loss aversion each day can be a method of ensuring that monthly payment obligations can be successfully managed, particularly if the work imposes disamenities (e.g. if it is boring, physically demanding, physically risky, etc.).[12]

From a firm's perspective where workers reaching a certain level of output matters most, it is advantageous to induce workers to engage in the strategic adoption of loss aversion if it leads to higher rates of goal attainment, which might be expected if present bias is common. How might firms do this? Firm strategies have generally taken two forms in the literature: 1) the use of loss-

worker will provide more effort on exogenously "good days" where the marginal costs of effort are low and less on exogenously "bad days" where the marginal effort costs are high, i.e. when $c_t'(e_t)$ is high. That is, when the marginal benefits are consistent from day to day, higher marginal costs will lead to lower equilibrium effort. Effort, therefore, will fluctuate from day to day. When there is present bias, beyond just a β discount, a worker has further incentive on "bad days" to implement effort substitution because of the expectation of future "good days" to make up for it. This is an implication of Proposition 2 in Koch and Nafziger (2016).

[11] "Daily targets can also serve a second purpose: like many mental accounts, they help mitigate self-control problems." (Camerer et al., 1997), pp. 426.

[12] This is directly supported by Dupas et al. (2020): both idiosyncratic stated income needs and expectations affect labor supply (their Appendix E). Stated daily income needs intensify reference dependence.

framed contracts with monetary incentives for output, where a reward for a minimum output is removed if the worker fails to produce a certain final output (e.g. Imas et al. (2017)) or intermediate/daily output (e.g. Kaur et al. (2015)); and 2) the self-imposition of short-run *mental* goals and narrow brackets that may or may not carry monetary incentives in order to achieve a longer-run target, as discussed above.

Imas et al. (2017) shows that not only do loss-framed contracts with discrete payments for total output lead to greater worker effort than gain-framed contracts, but workers actually *prefer* loss-framed contracts. The authors suggest that workers use the loss framing to keep themselves accountable dynamically to achieve a minimum output because they are otherwise dynamically inconsistent. Unfortunately, dynamic or period-specific behaviors are not observed in the study. However, one role of such contracts is clear: non-linear incentives establish expectations for worker final performance by codifying, in the framework above, a target $\sum_{t=1}^{T} b(e_0^*)$. There are two mechanisms at play: first, the bonus makes a final output target *salient*; second, it ties a lump-sum reward to it, thus raising the stakes of failing to stay on target and increasing the value of strategically self-imposing short-run loss aversion. The salience and monetary consequence of the target create expectations around which workers are more likely to plan and a target they would prefer to achieve ex ante, in the spirit of Kőszegi and Rabin (2006). The simplest form of narrow bracketing in this case is, therefore, to take the long-run bonus target and simply distribute it into daily targets, as in Koch and Nafziger (2020).

Kaur et al. (2015) show that when offered the chance to voluntarily decrease their piece rate if they did not achieve their chosen daily target, a large portion of workers (more than one-third) voluntarily chose this dominated contract as a way of solving self-control problems that day. This mechanically imposes loss utility. Unlike the use of long-run bonuses, these incentives reduced the piece rate for short-run output in a job setting that pays piece rates without longer-run performance targets and bonuses. Other than showing that workers prefer to impose commitment devices upon themselves, it is difficult to draw conclusions about the relationship between longer-run contracts (e.g. Imas et al. (2017)) and short-run behaviors. Because real-time performance is more costly to monitor and punish or reward in a variety of circumstances, many firms and industries use bonuses for longer-run output instead of tightly monitoring short-run output.

Thus, the prior literature leaves open significant questions: first, is there evidence that workers self-impose reference dependence and loss aversion in their daily labor supply choices (without loss-framed or dominated contracts), and are these employed to solve self-control problems?[13] Second, can firms, through their use of non-linear payments, establish long-run targets and expectations for workers and thereby *induce* workers to adopt narrow brackets in this way? More broadly, when a bonus payment is on the line (either loss- or gain-framed) and used as a longer-run

[13] In the sales data and experiment, I do not directly test for present bias but rather if workers behave and plan *as if they expect* present bias.

target, how do workers make decisions in the short run to generate greater long-run output, as they do in Imas et al. (2017)? How exactly do non-linear payments lead to greater performance (e.g. Kuhn and Yu (2021); Cai et al. (2022))? I empirically explore these dynamics through the use of real-world performance data from a door-to-door sales company and a real-effort task experiment. My analysis shows that firms do, indeed induce workers to adopt narrow brackets by establishing longer-run targets via a bonus scheme and that inducing them to do so increases their performance.

3 Door-to-Door Sales Context

The door-to-door sales industry constitutes a sizable portion of the "direct sales" industry. Workers in door-to-door sales are presented with high-powered incentives, including high commission rates that rise with performance and the use of bonuses. These are also common in a variety of sales occupations.

A large number of firms that engage in door-to-door sales are located in the Mountain West region of the United States and employ thousands of college-age workers each summer to sell their products and services.[14] These include solar panels, pest control services, knives, and home security systems. General industry practice is relatively homogeneous across these products. Recruited sellers meet with managers, listen to an explanation of the work and earnings potential, and sign independent contractor agreements that stipulate the commission structure under which they will sell and their assigned city. The work itself is unpleasantly hot in the summer and often entails distasteful interactions with local residents. To entice skilled sellers to join their teams under these conditions, most companies will advertise that sellers make an average of $40,000 during the late April to late August sales season selling six days a week. There is a high level of competition between companies seeking to land top talent, and there is an extremely wide variance in sales skills among recruits, leading to a large variance in income. The company whose data I analyze, which I will call "PestCo," operates within these norms.

A "sale" at PestCo is recorded when a resident signs a contract for pest control services that lasts 12–18 months for services given quarterly. The contract is recorded electronically. Within pest control sellers at PestCo, the timing of sales can vary widely. On average, sellers generate one sale for every 20 pitches they present, but exactly which of those 20 pitches will result in a sale and at what time each sale will occur is highly uncertain. Any single pitch could result in a sale, so each knock on a house door is akin to entering a type of lottery. Hitting one's expected number of sales early in the shift comes as a meaningful surprise. Because the value of each sale to the seller

[14] One reason for locating in this region is the large supply of young college students (usually age 20-25) who have recently returned from 2-year or 18-month proselytizing missions for The Church of Jesus Christ of Latter-day Saints, which is headquartered in Salt Lake City, and whose members are the majority in the state of Utah. These missions, in a purely practical sense, use skills very similar to a sales job: approaching strangers and striking up a conversation, connecting quickly, moving conversations toward a specific goal, and winsomely absorbing rejection. Recruiters understand this dynamic and seek to capitalize on these skills.

is large, the stakes for each sales pitch are high.[15]

PestCo, like nearly all door-to-door sales companies, pays large commissions in the range of 18–40% on the value of the service contracts they generate. A typical sale can result in an income to the seller between \$100 and \$250 depending on the value of the service contract signed by the customer and the seller's commission rate. Importantly, commission rates are increasing in cumulative sales performance and increase discretely in increments of 50 sales. The final commission percentage for each sale is calculated at the *end* of the sales season. The result is a discrete bonus with a small marginal increase in per-sale commissions past the threshold. Sellers are paid an up-front portion of their commissions (\$75 per sale) during two-week pay periods, similar to a regular paycheck. The balance of commission payments is calculated at the end of the season based on final performance and paid out thereafter.

Figure 1 characterizes a seller's total earnings at the end of the summer season depending on their total sales at an assumed contract value of \$500. A seller who produces 149 sales receives a commission of 25% on all sales at the end of the season, while a seller who generates 151 sales receives a 27% commission. This results in a lump-sum bonus of approximately \$1,500 for crossing the 150-sale threshold (plus approximately \$10 more per sale above it). In addition to this de facto bonus from the commission change, the seller receives a flat "rent bonus" of about \$2,000 that covers the seller's apartment rental costs for the summer. The average first-year seller yields between 100 and 150–175 sales, while experienced sellers generate 150 to 300 on average. The highest ability sellers generate over 350 sales for incomes in the \$60,000–\$80,000 range. Importantly, if a worker expects to end up in a particular interval, the main incentive is a piece rate with a bonus.[16]

The prior literature has generally considered occupations in which income is a smooth function of hours worked with relatively small deviations from average income. For example, a standard deviation in wages for a taxi driver is only about 10% of the mean (Thakral and Tô, 2021). At PestCo, a standard deviation in daily sales is 100% of the mean, and the effective daily wage can double in as little as 30–60 minutes. Income is accrued in discrete units, creating more salient opportunities for earnings references than in the past literature.[17] The "lumpy" nature of income in this context is an advantage over existing studies because each door interaction is quasi-experimental. The skill requirements of the job also make this setting unique in the literature. Sellers must be able to strike up a conversation with a stranger, understand and respond to objections, communicate the value of the product, and adapt their strategy on the fly as more information about the customer is revealed.

[15] Requests to come back to the house later yield similar success rates to knocking on a door for the first time. Callback requests are considered by sellers to be an indirect way of politely declining the sales offer.

[16] While the compensation is technically kinked at each threshold, the most important incentive is a flat bonus, not a kink, making this different from the recent literature. At 250 sales, sellers qualify for the company's all-expenses-paid vacation.

[17] E.g. it is much easier to count contracts sold than total income earned net of tips while driving, even when the tips are "large" (approximately \$30) as in Schmidt (2018).

Each of these tasks is cognitively demanding and requires strong interpersonal skills.

Another unique feature of this setting is that outside considerations that might influence the formation of medium- and long-term earnings targets in other settings are absent from this setting. Most sellers are below the age of 25 and have not formed financial commitments that require set payments that might influence the formation of salient short-run "income needs" as examined in a prior study (Dupas et al., 2020) or long-run payment obligations. The cost of housing, for example, is paid for up front by PestCo and repayment is not required unless the seller fails to secure 150 sales. These needs are, therefore, baked into the performance schedule. Income needs over the short-run (e.g. the week) or over the medium-run (e.g. the month) are not fully operable because the vast majority of compensation for their work is received by workers at the *end* of the season. Similarly, because these workers live away from their normal homes and networks, there is limited scope for daily external obligations to shape their daily labor supply choices.

Through the company's internal website and mobile app, sellers can view their performance history. All workers are aware of their normal performance, including their cumulative sales and average output each day. The availability of this information makes references related to one's own performance highly salient. Through its website and mobile app, PestCo tracks every sale and house "knock" recorded by each seller. This forms the basis of my analysis dataset. See Appendix D for more details on industry practice and contracts.

4 Data

My analysis datasets come from the comprehensive sales and seller tracking databases from PestCo for 2018–2019. The company uses a common sales tracking app that documents every door at which a seller records interacting with a resident and the location and timestamp of those interactions. PestCo separately tracks the date and time each service contract is signed, the location of each customer, and the seller who generated the sale. Together, these two systems give a comprehensive view of the activities of each seller every day they are knocking on doors and selling in their work area.

Using the raw sales and knocking data, I construct two panels of individual seller performance. First, I build a daily panel of each seller's sales, work hours (defined as the time between the first knock/sale and the last knock/sale), cumulative sales over the season, and cumulative average daily sales as a measure of "recent expectations." Following the past literature (Crawford and Meng, 2011), I calculate a proxy for each seller's recent expectations by examining each seller's average past daily sales during the season. The selling week runs Monday through Saturday, and because residents are home at higher rates on Fridays and Saturdays and seller experiences differ by day of the week, I calculate each seller's average daily sales specific to each day of the week from all past days in the same sales season. These expectations can update and evolve over the course of a season, though the measure is remarkably stable after the first 2-3 weeks.[18]

In my second dataset, I construct a panel of each seller's pitches presented to a prospective customer, daily cumulative sales, and stopping probability each half hour of their shift. This interval of observation is the same as that in the recent taxi literature (Thakral and Tô, 2021). For each seller in each half hour, I create a measure of their current distance to their daily expectations: their number of cumulative sales so far that day minus their average sales for that day of the week. For values less than zero, a seller has not yet achieved her expectations and is therefore in a loss domain, while values greater than or equal to zero indicate a seller is in a gain domain. In this dataset, I define "starting" a shift as the half hour of the day in which a seller records her first knock of the day, and I define "stopping" as the half hour of the shift when the last knock of the day was recorded. In all, my half-hourly panel contains approximately 459,000 observations for 512 sellers across 180 days in 2018-2019 covering the late-April to mid-August season.

I supplement these panels with daily weather data from the National Oceanic and Atmospheric Association (NOAA) National Climate Data Center (Menne et al., 2012). I include daily total precipitation, high temperature, and low temperature from the weather station nearest to each seller's working ZIP code as controls. These factors may be important because door-to-door sales is an almost exclusively outdoor job. During these summer months, heavy rain and humid heat greatly increase the marginal cost of effort, and heat can have negative effects on cognitive ability and learning (Park et al., 2020). Alternatively, these factors might keep people inside their homes if the outdoor conditions are inhospitable, so the relationship between sales and these conditions is ambiguous ex ante.

One theoretical concern in this context is that sellers might be differentially sorted by managers into neighborhoods that are "easier" or "harder" to sell in. I include in my analysis controls for the characteristics of each person's work area. I use ZIP code data from the American Community Survey's 5-year summary files for 2013-2017 to serve as controls. I include variables that are likely to affect demand for pest control services or the ability to pay for them.[19] However, there is essentially no evidence of sorting behaviors correlated with seller performance. Additionally, managers emphasize that making assignments to work areas based on perceived skill or other seller attributes is costly to them as managers and generates unclear returns, which significantly undermines the business case for it.

[18]Using various definitions of recent expectations such as sales in the prior five weeks yields similar results. I also use performance from the first two weeks of the season to predict forward the seller's nearest "goal-based" reference, which yields similar effects. See Table A4.

[19]These variables are median household income, rates of unemployment and poverty, the share of home values in specific ranges, total housing units, the share of units that are owner-occupied or are single-family homes, the share of households with a married couple, the share of adults with a Bachelor's degree or more, and the share of the population that has not moved in the past year. These jointly explain less than 3% of the variation in sales both between and within sellers. Appendix Table A2 provides details of a regression of daily sales on weather and ZIP code characteristics. Only three coefficients are statistically significant at the 10% level. Estimates excluding these controls are nearly identical but slightly less precise. See Appendix D for additional background.

Summary statistics for my two panels are in Table A1. Across all half-hour periods, the average number of sales is 0.16 based on 2.28 pitches. The average number of sales per day across all sellers is approximately two based on 6.9 hours per day, though there is substantial variation. Sellers work in relatively high-income areas. The median household income in their sales areas is $86,000, and nearly 20% of residents in the average ZIP code have incomes between $100,000 and $150,000. Seller work areas are mostly single-family homes (mean of 80%), are predominantly non-Hispanic white (mean of 80%), are relatively highly educated (mean of 45% Bachelor's degree or more), and have stable populations.

From the half-hourly panel, Figure A1 shows the distribution of start and stop characteristics for each working day. Panel A shows that most sellers start their shift with their first knocks and sales between 1:00 PM and 2:30 PM, though there is substantial variation in start times. Some start as early as 10:00 AM, while others begin working in the late afternoon or early evening. After starting their shift, the majority of sellers stop working between the sixth and eighth hours, though a large share stop working for the day before their sixth hour of work. This context and the availability of comprehensive data provide a unique opportunity to test for real-world reference dependence connected to firm-imposed incentives.

4.1 Sales Context and Theory

What does the theoretical framework in Section 2 imply for my sales setting? Here, a descriptive example is helpful. When PestCo sets a bonus at 200 sales, the bonus directly affects a forward-looking worker who knows her ability on the job could reasonably yield her something close to that number (e.g. 190 sales). In the case of both the standard model and the KR model, she may raise her objective for total sales at the end of the season to be at least 200 (a new $b(e_0^*)$ in the framework of Koch and Nafziger (2020)) because she believes it is attainable and the $2,000 bonus is worth the extra effort.

Next, if she expects present bias may impede her from achieving her 200-sale target, she may engage in narrow-bracketing. Knowing she needs 200 sales over 100 days, she can set expectations for each day's performance: just over two sales per day. She then works with these two sales per day as her reference point, which satisfies the personal equilibrium condition. Being below her two sales generates negative comparisons and a higher marginal utility for each sale, so she will work harder or extra hours to get the remaining sales. If she does achieve her two sales, she can then quit for the day and feel satisfied with her performance as her marginal utility has declined, leading to a kink in labor supply at expectations. Achieving her two sales then keeps her on track to hit her goal of 200 by the end of the season. In the absence of the 200-sale bonus and a single piece rate, she may have maintained her trajectory of lower total performance (190 sales) and would not have engaged in narrow bracketing, so there would not be a structural break in daily labor supply around expectations.

An important prerequisite for narrow bracketing behavior being amenable to goal-setting is that the worker must have a realistic, forward-looking view of what she can plausibly accomplish, i.e. not full naivete about her level of present-bias (β) and her effort costs. In other words, like the above example, she would need to reasonably predict her ability to achieve approximately 190 sales. In my sales setting, a sign that sellers are forward-looking would be that their daily labor supply does not substantially change as their cumulative performance (and therefore realized commission rate) increases because they have already optimized for their long-run expectations. In other words, she would not respond to reaching 100 or 150 sales and yielding an increase in her commission rate because she already expected to reach a total above 150.[20]

The prior empirical literature on reference dependence has been unable to examine short-run goal-setting as a response to a long-run target because the work settings analyzed to date do not provide a clear endpoint at which a worker evaluates any long-run goals she may have. The "long run" is too nebulous. On the contrary, my sales setting provides a clear end date. A second reason the prior field literature has been unable to examine long-run goals is that the occupations under study are measured in settings in which other factors such as income needs may form the most salient (or only) form of medium- to long-run targets, which remain unobserved to the researcher.[21] PestCo, through its use of lump-sum bonuses, provides external incentives for workers to set their sights upon specific long-run outcomes just as in my experimental setting. These bonuses increase the salience of particular points to act as targets and attach a significant monetary incentive to achieve it. Meanwhile, the sales setting is a fixed-term job that is conducted far away from "home" and is paid mostly at the end of the season among young workers without major fixed-schedule financial obligations whose housing costs are baked into the bonus schedule. This limits the scope for outside income needs to dictate specific points in the earnings distribution as targets. Thus, this empirical setting provides a unique opportunity to study these questions empirically.

4.1.1 Theoretical Predictions

Based on the theory discussed previously, if daily reference dependence and goal setting occur in this setting, I expect:

(A) There will be a kink and/or discontinuity in labor supply upon surpassing daily performance expectations. No such shift is predicted by the standard model.[22]

[20] While this assumption of a non-myopic view is reasonable, it is not certain; evidence from other contexts indicates that myopia affects the optimality of decision-making in areas like pension planning (Mitchell, 1988), health behaviors (Cawley and Ruhm, 2011), and take-up of financial aid (Bettinger et al., 2012).

[21] Bénabou and Tirole (2004) propose that self-reputation is what animates the use of daily targets. However, they do not consider the use of longer-run financial targets or needs in their models. In a variety of economic interactions, reputational considerations are intertwined with economic incentives such as promotions, bonuses, or the option value of future job prospects.

[22] Likewise if the worker exhibits loss aversion around expectations only over *hours* rather than sales *performance*, conditional on hours worked during the day, there should be no kinked relationship between sales performance and work exertion.

(B) As a result of the change in marginal utility demonstrated by (A), the distribution of performance around expectations should be narrow and subject to bunching relative to a counterfactual in which reference dependence is absent (or less in degree). If these daily references are connected to the bonus scheme, these distributions should form around linear "paths" to reach the relevant bonus as in the experiment, i.e. the average each seller would need each day to reach the bonus.

(C) If workers have established their daily targets *as a commitment device* to achieving the long-run objective (the bonus), those that have surpassed their relevant bonus will reduce their effort even though they are still paid a significant piece rate for each sale. The variance of effort each day in the final days of the season should increase, while the variance of cumulative performance should decrease because workers above the bonus lower their effort while those below maintain their effort. This rules out status quo anchoring or habituation as explanations. Reducing labor supply after achieving the target, by itself, does not require a behavioral mechanism, but it does indicate that these workers are not simply anchoring their reference dependence to *expectations* separate from their *goals*—in other words, the two are directly connected.

These predictions are similar in the experiment setting I present in Section 7.

5 Empirical Strategy

5.1 Tests of Reference-Dependent Labor Supply

I first use my half-hourly panel to test for the presence of reference dependence in daily labor supply choices consistent with Prediction (A). As outcomes, I focus on stopping work for the day, a measure common to the past literature, as well as pitches presented in the next half hour, a measure of effort "exertion" conditional on continuing to work.[23]

My empirical approach approximates an experimental ideal in which sales performance is randomly assigned each half hour by netting out conditions correlated with effort costs and the number of sales a seller has generated to that point. The underlying assumption is that conditional on my various fixed effects and controls, the exact number of sales a seller has at a particular point in the day is as good as random. Given the context in Section 3 and the set of controls and fixed effects I present, this assumption is reasonable. The sales setting presents a unique opportunity to study this behavior because whether a sale occurs or not depends strongly on who answers the door when a seller knocks—similar to a small lottery.

I first estimate a non-parametric model of labor supply with respect to each seller's distance from their sales target to trace out patterns without imposing a functional form. Following the past literature (Crawford and Meng, 2011), I define expectations and targets in all my models as

[23] In Appendix A, I also examine the probability of recording any knocks during the next half hour. This measures whether workers are more likely to take a break as a result of their position relative to expectations. These results closely mirror the knocks-based exertion margin. See Appendix Figure A2.

the average daily sales from all past workdays in the season specific to each day of the week (i.e. a specific mean for Mondays, Tuesdays, etc). I note here that PestCo runs various competitive tournaments during the sales season of three different types. Because these significantly change the incentives faced by the sellers and may shift the workers' target for the day, I separately analyze behavior during non-tournament days and present those results in my tables and figures.[24]

For seller i during half hour of the shift t and half hour of the day h on day of the week d in week of the season w in year a, I estimate the following model:

$$\begin{aligned} y_{ithdwa} =& \beta_0 + \sum_{e=-k, e\neq 0}^{k} \beta_e * \boldsymbol{I}_e\{sales_{ithdwa} - \overline{Sales}_{idwa} = e\} \\ &+ \alpha X_z + \sigma W_{dwa} + \mu_{it} + \eta_h + \nu_d + \omega_w + \tau_a + \varepsilon_{ithdwa} \end{aligned} \tag{1}$$

Here, y is the probability of stopping work for the day as well as the number of pitches presented to a resident in the next half hour. The expression $\{sales_{ithdwa} - \overline{Sales}_{idwa} = e\}$ represents the seller's current distance to expectations: her current cumulative sales that day ($sales_{ithdwa}$) minus the worker's average daily sales specific to the day of the week ($\overline{Sales}_{idwa}$). $\boldsymbol{I_e}$ is a dummy variable assigned to each distance value. The coefficients of interest, β_e, capture non-parametric effects of being e distance from one's expectations target. Distance values below zero are characterized as being "losses" and values above zero are "gains." Because sales are discrete values, these coefficients include values rounded to the nearest integer, with the (0,1.5) interval being included in β_1.[25] Under reference dependence, beginning with β_1 there will be an upward change in stopping probability or a downward change in exertion as the distance from expectations increases.

The various fixed effects (μ_{it}, η_h, ν_d, ω_w, τ_a) are for seller by half hour of the shift, half hour of the day, day of the week, week of the season, and year, respectively. Importantly, μ_{it} captures a seller-specific baseline hazard over the shift. That this factor is omitted by the prior literature is noted by Thakral and Tô (2021). They include a driver-specific hazard in their estimates of taxi driver behavior and conclude this is vital for unbiased estimates of labor supply responses to daily earnings. I incorporate this methodological improvement into my estimates. The X vector is the set of ZIP code characteristics from the ACS, and W is the set of weather controls from NOAA discussed in Section 4. Importantly, these controls rule out any relationship that might arise between sales and factors correlated with the length of the shift (e.g. running up against a

[24] I also estimate my parametric models using a pooled sample across all tournament and non-tournament periods and interact my coefficients of interest with indicators for tournament periods. These estimates are in Appendix Table A4. I report the non-tournament coefficients in my figures (see Section 6.1.1). See Appendix D for more on these tournaments.

[25] Other studies examining reference dependence discretize earnings into ranges. The "correct" size of the earnings range has been the topic of some disagreement (Farber, 2015; Martin, 2017; Thakral and Tô, 2021). In sales, earnings are already discrete, so I do not have to impose a bin structure. Because the common support in the distance to expectations is thin outside the [-4,4] interval, I plot that interval in my figures. I report the full set of distance dummy coefficients corresponding to my figures in Appendix Table A3.

maximum work hours limit or fatigue), the time of the day (e.g. an 8:30-9:00 PM hard stop time), weather, or cross-sectional differences between each worker in their "normal" work schedule.

In my main models of interest, I fit parametric estimates that impose a functional form to match the non-parametric estimates in Equation 1, with linear splines divided at zero:

$$\begin{aligned} y_{ithdwa} =& \beta_0 + \beta_1\{sales_{ithdwa} - \overline{Sales}_{idwa}\} \\ &+ \beta_2\{sales_{ithdwa} - \overline{Sales}_{idwa}\} * \boldsymbol{I}_{sales \geq \overline{Sales}} \\ &+ \beta_3 * \boldsymbol{I}_{sales \geq \overline{Sales}} \\ &+ \alpha X_z + \sigma W_{dwa} + \mu_{it} + \eta_h + \nu_d + \omega_w + \tau_a + \varepsilon_{ithdwa} \end{aligned} \tag{2}$$

This approach allows the slope of the relationship between labor supply and distance to one's reference point to differ in the gain and loss domains. $\boldsymbol{I}_{sales \geq \overline{Sales}}$ is a dummy for if current sales are above expectations, or in other words, for entering the gain domain. β_1 defines the slope of the relationship between one's current distance to average sales and labor supply in the loss domain. β_2 captures the change in slope upon crossing the reference and entering the gain domain. Finally, β_3 captures any discontinuous level shift in stopping probability or effort from reaching the reference. The fixed effects and controls are all the same as in Equation 1. In a standard framework, there should be no sudden change in the slope and no discrete level shift upon reaching the reference point. Under reference dependence with loss aversion, we would expect to see an upward change in the slope of stopping probability. In other words, β_2 will be significantly positive in the stopping model. The coefficient β_3, while not predicted by simple loss aversion, represents a discrete penalty for "losing," or for falling short of expectations, which suggests reference dependence.[26] If β_2 and/or β_3 are significant and positive in the stopping model, this represents strong evidence of reference-dependence.

5.2 Tests of Goal-Setting and the Bonus Schedule

I next use my daily panel of sales and labor supply to examine goal setting by sellers around the bonus schedule. First, to visualize how workers respond to bonus incentives with the predictions discussed above, I present kernel density estimates of cumulative sales at the end of the season and as well as throughout the season for all workers. I present densities at two-week intervals to illustrate the evolution of sales over time. I also perform the same analysis for subgroups in particular total performance bins from the end of the season to trace how the densities within groups progress (relevant to Prediction (B)). As the focal example, I present these for those whose total sales at the end of the season were between 175 sales and 225 sales, putting them around the bonus threshold of 200 sales. If workers with the same apparent goal at the end of the season have a narrow and/or bunched distribution of performance, this further suggests that workers are

[26]Estimating with second-order polynomials results in small and statistically insignificant coefficients on the squared term for both outcomes.

anchoring to their goals and exhibiting more effort while below their daily expectations, which compresses the distribution upwards.

A prerequisite for setting long-run goals is that sellers must not be myopic. To test for this, I estimate how sellers adjust their labor supply as their cumulative sales increase throughout the season. Sellers only know their final earnings per sale at the *end* of the sales season after their total number of sales and total revenue are calculated. If sellers have realistic, forward-looking expectations for what they can achieve (perhaps after an initial learning period of a few weeks), perceived changes in their wages that come with entering a new 50-sale performance interval should not change their daily labor supply because they have already optimized over their chosen long-run outcome. According to the KR model and the standard model, those with higher expectations for their commission rates should work more hours than those with lower ex-ante expectations. Conversely, myopic agents would respond to an increase in their realized commission rate, which is inconsistent with long-run planning. In Appendix B, I show that daily performance across the season is remarkably stable: mean performance after the first two weeks of the season explains 75% of the variation in final season sales, while sales after the first 5 weeks explain nearly 90%. This also indicates very little switching of targets and a consistent personal equilibrium after the first few weeks. Sellers appear to select bonus thresholds early on and work consistently with those targets each day.

According to Prediction (C), sellers that have previously worked in the pursuit of their long-run targets should reduce their labor supply once they surpass (or will imminently surpass) their long-run targets. While a reduction in labor supply could occur with a standard agent, the combination of daily loss aversion and reduced labor supply upon nearing the bonus suggests a direct connection between the worker's current expectations (i.e., their personal equilibrium) and their goal to achieve the bonus.

To test both of these dynamics, I use my daily panel to regress hours worked per day on indicators for 10-sale intervals of current cumulative sales interacted with indicators for 50-sale bins of total sales at the end of the summer. I estimate the following equation for seller i on day of the week d in week of the season w in year a working ZIP code z:

$$\begin{aligned} y_{idwa} =& \beta_0 + \sum_{e=[0,10)}^{[320,330)} \sum_{f=[100,125)}^{[300,325)} \beta_{ef} \boldsymbol{I}_e * \boldsymbol{I}_f \\ &+ Efficiency_{idwa} + \alpha X_{zdwa} + \sigma W_{zdwa} + \mu_i + \nu_d + \omega_w + \tau_a + \varepsilon_{idwa} \end{aligned} \tag{3}$$

The outcome variable y is the number of hours worked per day. The indicators $\boldsymbol{I_e}$ and $\boldsymbol{I_f}$ are indicator variables for currently working in interval e and for having total season sales in interval f. In this specification, β_{ef} captures the non-parametric effects of being in interval e for an individual whose total sales for the season were in interval f. These coefficients trace the labor supply path of those who ended with a similar total number of sales. The X and W vectors are the same as Equation 1. The $Efficiency$ variable is a time-varying measure of each seller's average sales per hour for all past workdays that season, which proxies for sales ability and may evolve as the season progresses. Changes in this measure capture learning effects over the season, which shifts the expected marginal earnings of an additional period of work. I include fixed effects for seller (μ_i), day of the week (ν_d), week of the season (ω_w), and year (τ_a). These fixed effects ensure that the β coefficients characterize within-seller choices holding constant other characteristics of the sales season, fatigue, or learning effects. If the β_e coefficients are constant within different types of sellers f as they cross intermediate 50-sale intervals, then it does not appear that sellers are responsive to a change in their realized wage.

The KR and standard models predict that the coefficients on all intervals in e should be consistently larger as their expected total sales—therefore, expected commissions—increases in f. Importantly, if sellers are focused on reaching a bonus threshold, the coefficients for β_e will be much smaller after crossing the worker's final bonus threshold. This would result in a significant drop in hours worked. For example, a worker who finished with 150-175 sales (just beyond the 150-sale bonus threshold) would work fewer hours in the intervals just at or after the bonus threshold. Equation 3 captures this dynamic for each bonus threshold from 100 to 300 sales.

6 Sales Results

6.1 Do Workers Exhibit Reference-Dependent Labor Supply?

Figure 2 shows each of the coefficients from the non-parametric estimates from Equation 1 as well as the linear estimates from Equation 2. Panel A indicates that as sellers approach their target from the loss region, the probability that they stop working for the day is relatively flat at a slope of 0.0021. After surpassing their target number of expected sales, there is a clear upward kink in the probability of stopping work. The slope of the relationship between cumulative sales and stopping probability in the gain region for expectations-based targets is 0.0058, or 2.8 times that in the loss region. For context, the average probability of stopping right at the reference point is 0.079, so an increase in this probability of 0.0058 for each sale past the reference point represents a 7.3% increase.

Panel B of Figure 2 shows the same estimates for pitches during the next half hour, which is a measure of exertion conditional on continuing to work. In contrast to the results for stopping probability, there is a relatively smooth relationship between exertion and sales each day and minimal change in this measure at the reference point. The size of the decline is small in percentage terms: each sale reduces effort conditional on continuing to work by approximately 1% from a baseline mean of 2.38 pitches. These results suggest that reference dependence in exertion is negligible, but that there is a steady decline in effort as sales increase. Reference dependence is most apparent at the extensive margin. In other words, if sellers stay on the job after reaching their expectations, their exerted effort is similar.[27]

Next, I use my estimates to calculate the parameter of loss aversion, λ. My setting requires an approach to measuring loss aversion that is not dependent on the measurement scale of the output units (sales). One such approach is advocated by Köbberling and Wakker (2005). Their measure focuses on the difference in the slopes of the utility function in the gain domain and the loss domain. Because my empirical model partials out all covariates correlated with effort costs and because the timing of sales is conditionally random, the only difference between the gain and loss domains is the difference in the marginal benefit, i.e. the ratio of slopes for each outcome measures $U'(0)_{\uparrow}/U'(0)_{\downarrow}$. The ratio of slopes in the stopping model is $0.0058/0.0021$, or 2.8, and the slopes at the "exertion" margin have a ratio of 1.5. Using a bootstrap with 250 replications for inference on the ratio of these slopes in the stopping model yields 95$\%$ confidence intervals for my estimates of loss aversion of 1.2 to 4.4 (1.5 to 5.3 for the bias-corrected interval). My estimate of 2.8 in my baseline models is the most conservative of my stopping model estimates, and my non-parametric estimates imply an even larger ratio. Other specifications, which I detail in Section 6.1.1 yield estimates as high as 4.1 or 5 for the stopping model and 3.9 at the margin of effort conditional on continuing. In their review, Gächter et al. (2007) find loss aversion of approximately 1.4 to 4.8 across various measurements, with an average value of 2.6. A coefficient of loss aversion in my results of 2.8 is, therefore, very consistent with the prior literature.

One notable feature of the KR model is that personal equilibrium is established immediately following the formation of expectations, even when the final payoff is far away. If reference dependence with loss aversion around long-run goals is operating in daily labor supply choices, then there should be evidence of this phenomenon across time periods, even early in the sales season when exact performance relative to the bonus threshold is not immediately payoff relevant. To investigate this, I separately estimate my models during the May, June, and July months. In May, opportunities for effort substitution are more plentiful, and these opportunities fall during June and July. Appendix Figure A6 presents these estimates. The evidence is consistent with my overall results, even in May and June. Similarly, if expectations are stable throughout the year, we might expect similar behaviors from sellers who have experience and those who are new to the job. Panel D shows that this is the case. There is little difference in behavior around expectations for those with more experience with the possible exception of experienced workers being slightly *more* likely to stop working for the day right at the zero cutoff.

[27] In two of my later specifications, there is a downward kink and a discontinuity in effort. In Appendix A, I also examine the probability of actively knocking, meaning recording any knocks during the next half hour. These results in Appendix Figure A2 closely mirror the exertion margin. At a mean active knocking share of 80%, the slope estimates are quantitatively small and not economically meaningful. Sellers do not appear more likely to take breaks during their work as a function of their position relative to expectations.

6.1.1 Robustness and Alternative Specifications

Rather than separately estimating stopping behaviors for non-tournament periods, my first alternative specification pools together all tournament and non-tournament periods and interacts each of my key measures of distance to daily expectations with indicators for what kind of tournament or non-tournament is operating. This allows the effect of crossing the reference to differ based on period type. The results of this specification are in Figure A7. The result for non-tournament periods is a more pronounced upward kink in stopping probability and the emergence of a downward kink in pitches per half hour. The slope in the gain domain is 4.1 times that in the loss domain for stopping probability, meaning loss aversion in this model is higher than in my baseline model. At the exertion margin, the slope in the gain domain is 3.9 times that in the loss domain. This specification confirms the results of my baseline model and provides even stronger evidence for reference dependence.

My estimates impose a linear structure with a cutoff at each seller's cumulative average sales. This choice is in line with the KR model of reference dependence around recent expectations. As a robustness check against incorrect specifications of the cutoff at zero, I estimate my models again using non-linear least squares. To incorporate my fixed effects and controls into my specification, I first residualize the probability of stopping with my full battery of fixed effects and controls and use the residuals in my non-linear least squares estimate. Rather than imposing slope and intercept coefficients at zero, I allow the cutoff itself to be a parameter of the model. The results are in Table A5. The non-linear least squares estimates confirm that there is, indeed, a structural break at the worker's average cumulative sales and a strong upward tilt in stopping probability. The exact cutoff in the non-linear least squares estimate is 0.11, approximately one-tenth of a sale from my measure of the seller's expectations, which may be consistent with the "buffer" idea against later fatigue, which is explicitly invoked by respondents in my online experiment.[28] The ratio of slopes in the stopping model is 5.1, meaning that my baseline estimates may be quite conservative. For pitches per half hour, even though the estimates show a statistically significant kink downward and that the ratio of the slopes across the reference is 3.3, the magnitude is small in percentage terms; each sale past the reference leads to a 1.8% decline in effort conditional on continuing to work, and the results are more sensitive to specification.

As an additional test, I estimate my baseline model but include exertion effort on the right-hand side: cumulative pitches that day as a measure of total exerted effort. If a worker is exerting a high level of effort on the job and becomes fatigued, the fatigue could be affecting her willingness to continue working or to exert effort in the next half hour. Table A6 presents these estimates for my parametric models. The results are nearly identical to my baseline model. The results for stopping

[28]This is consistent with Kőszegi and Rabin (2007), who build on their concept of personal equilibrium in their theory of "preferred personal equilibrium." Here, the anticipation of risk leads to a strong tendency toward planning and the purchase of insurance. This performance buffer may, therefore, be a form of small-scale insurance.

behavior imply that my baseline model adequately controls for effort differences at the intensive margin that may have generated differences in sales. At the exertion margin, the negative slope in the loss domain is not as steep as my baseline model. Upon entering the gain domain, there is essentially no change in the slope from the loss domain, indicating that the decline in pitches across the reference is smooth.

I next create an alternative measure of each worker's daily reference point and also estimate my models with the full tournament/non-tournament interaction. I construct a "goal-based" reference by examining the first 2 weeks of the worker's performance. I project their average daily sales from this period to the end of the season and then round to the nearest bonus threshold. If workers are projected to be within 15 sales of a bonus, I round up to the bonus, but if they are less than 35 sales over a bonus, I round down. I base this on the pattern of bunching from the kernel density estimates in the next section (6.2). I then allocate the average daily sales the worker would need to achieve this nearest bonus. These "goal-based" references are highly correlated with my proxy for recent expectations (0.82), consistent with a worker's rational expectations matching her likely goals. In Appendix Figure A8, I show that the use of this reference point is consistent with my baseline results.

6.2 Is the Bonus Schedule Inducing Reference Dependence?

I now present the distribution of sales throughout the sales season, which is relevant to predictions (B) and (C). Figure A3 shows the results of kernel density estimates for total sales at the end of the season. Around each 50-sale bonus threshold, there is significant bunching, particularly at 150 and 250 sales when the bonuses include rent payments and the company vacation. This indicates that the bonuses are salient for the sellers. Figure 3 breaks down the density of total cumulative sales for each seller in two-week increments over the season.[29] Unevenness in the estimated density graphs is apparent beginning in week four and becomes clearer in weeks 8–10, which is just over the halfway point in the season. Notably, bunching groups that form early persist further up the sales distribution over time.

An even starker pattern emerges when examining groups of workers with a similar total performance at the end of the season. Figure 4 presents the kernel density estimates of cumulative sales over the same two-week intervals as Figure 3, but I limit this to those whose total sales at the end of the season were between 175 sales and 225 sales, or those around the bonus at 200 sales. In week 4, the distribution is tightly centered, after which bunching emerges in the distribution. This persists until approximately week 12, at which time the growth rate of the top of the distribution starts to slow, which compresses the distribution from the top. These figures confirm that sellers are particularly cognizant of and responsive to these lump-sum bonuses.[30] Visual evidence of upward

[29]If a seller left relatively early in the season, their sales are included in the total as of the date they left and hold the same value as the weeks progress, so the relatively high density below 100 includes those who only worked a portion of the season.

[30]A similar pattern is visible for those who finish the season around the 150-sale bonus threshold, as seen in

pressure from the left tail of the distribution, especially in the early to middle stages of the season, is consistent with predictions for workers with reference dependence. Returning to the results from the half-hourly panel, it is important to reiterate that if this pattern of tight performance were based on workers anchoring their day to particular work hours, then there would be no detectable change in the relationship between stopping probability and distance to sales expectations conditional on the length of time worked and/or time of the day; I find such changes in Figure 2.

I next show evidence that sellers are forward-looking and that they significantly reduce their labor supply upon nearing or passing their relevant bonus threshold. Estimates from Equation 3 are summarized in Figure 5, which shows the predicted hours worked per day over 10-sales increments of cumulative sales from this model. Each line shows the labor supply trajectory of different bins of total sales at the end of the season. Sellers do not appear to be myopic. After an initial adjustment period, those whose sales totaled over 300 quickly began working approximately 8 hours per day, while those with fewer than 200 total sales worked approximately 7 hours per day consistently over their accumulated sales. This difference is consistent with the both the KR and standard model's prediction that *expected* increases in wages would increase labor supply. Notably, within tiers of total sales, there is very little variation in the predicted hours worked each day over current cumulative sales, and labor supply does not significantly shift upon receiving a commission raise by crossing into a new 50-sale interval. These results show that workers do not change their work hours regardless of how much of a commission increase they have secured, suggesting a singular focus on long-run performance expectations.

These patterns shift significantly once the relevant bonus threshold has been reached. Even when conditioning on elapsed time in the season, weather patterns, and efficiency gains, sellers drastically reduce their work hours by 1.5 to 2.5 hours per day (20-30%) after passing the bonus.

Given the above observations, it is useful to examine the variance of performance each day (an exercise I also perform in the experiment below). I examine this within-day effort variance by looking at the standard deviation of worker daily hours across two dimensions: time and distance to each worker's final sales tally. To remove any composition effects that might drive this variance on any particular day, I regress each seller's daily work hours on a set of seller and day-of-the-week fixed effects. This removes volatility attributable to day-of-the-week effects and worker composition effects, i.e., who shows up to work that day. I present the standard deviation of the distribution of the subsequent residuals by elapsed time in each seller's season and by distance to the seller's final total for the summer.[31] In both instances, the variance of the sellers' daily effort should be consistent until the very end of the contract, either in the final days or in the final sales. A sharp rise in the variance of labor supply signals a departure from the worker's prior behavior,

Appendix Figure A4.

[31] When examining distance to one's total tally, I also include a fixed effect for week of the season to distinguish progress toward the goal from time effects.

consistent with the bonus being the impetus behind the initial daily loss aversion.

Figure A5 shows that the standard deviation of residualized seller labor supply remains relatively flat over the course of time and distance to the seller's total until the very end of the contract. In Panel A, while the variance of hours decreases a small amount during the first month as new sellers learn about their capability in the field, there is very little change from day 30 until approximately day 85. Panel B shows that the variance of seller labor each day does not systematically vary until the goal or final tally is within approximately 50 sales.

These aggregate patterns may mask composition effects; for example, those that achieved 100 sales in total are not represented in Panel B when the distance to the final tally is more than 100 sales away. To investigate this further, Panels C and D show the same phenomenon as Panels A and B but are separated by groups of total sales (0-100, 100-200, 200-300). Panel C shows that the rise in the variance of seller labor supply is most concentrated among sellers with over 200 sales during the season, though there is a rise in the variance among those below 100 approximately halfway through the sales season. These workers are more experienced, on average, and so are likely to have a better sense of when or how they may reduce their efforts around their bonus threshold. Panel D shows that all these groups exhibit an increase in the variance of their labor supply in the final 25-50 sales.

Taken together, the results of each of these exercises in Sections 6.1 and 6.2 show that these sellers are 1) able to predict their own performance very early in the sales season; 2) aware of and responsive to the bonus schedule; 3) setting goals around bonus thresholds in the schedule; 4) distributing their long-run goals into daily expectations; and 5) shifting out of their prior patterns (i.e. their loss averse daily labor supply) upon reaching or surpassing the bonuses. That all five of these hold empirically is consistent with the use of narrow goals in pursuit of long-run objectives (Shefrin and Thaler, 1992; Camerer et al., 1997; Koch and Nafziger, 2016; 2020) and that the firm's bonus schedule induced the behavior. An added strength of the online experiment, which I explain below, is that I can directly and causally attribute short-run reference dependence to non-linear compensation schemes in a similar compensation context.

7 Experimental Design

As an additional test of these dynamics in a controlled setting, I conduct an online real-effort task experiment on the Prolific platform. The participants engage in a simple button-pushing task in which they are asked to alternate pressing the "a" and "b" buttons on their computer keyboard, following closely the procedure in DellaVigna and Pope (2017). A successful sequence of "a" and "b" results in 1 point. Participants were asked to perform the task for a total of ten minutes in four rounds lasting two minutes and thirty seconds for each round with a break of ten seconds between rounds.[32] A total of 1,464 recruits completed the task.

[32]Important for this setting is that dynamic inconsistency has been shown to occur over time periods measured in minutes (McClure et al., 2004; Brown et al., 2009).

Each participant was paid a flat $3 payment to participate. I then randomly presented participants with one of three incentive conditions:

1. A bonus of $1 for achieving 2,000 points (2,000 bonus condition)
2. A piece rate of $0.05 per 100 points (the piece rate condition)
3. A bonus of $1 for achieving 2,400 points (2,400 bonus condition)

Importantly, the payment rates were calibrated based on the distributions of performance in DellaVigna and Pope (2017) to have equal predicted mean performance over the full ten minutes, meaning the expected payoff for a performance of 2,000 points (the approximate mean performance in that experiment) is exactly equal in the first bonus condition and the piece rate condition. The core difference between the two is that the bonus makes the 2,000-point target salient and payoff-relevant for the end of the 10-minute task period. The bonus environment in my sales setting follows this setup closely.

The Koch and Nafziger (2020) model predicts that if participants are induced to set a target by the bonus offer, optimality suggests they will set narrow brackets for themselves with their target being some point at or above the average number of points they would need to achieve to reach their target by the end of the task. They would exhibit loss aversion around this target as a commitment device in this tedious task, even in the present of little or no present bias. As mentioned previously, even in the case of the self-employed or pure piece rates (like in the taxi driver case), workers may exhibit reference dependence in pursuit of a longer-run target. The imposition of goals by, for example, a firm can *intensify* the use of internalized loss aversion as a commitment device to increase performance, particularly if there is present bias. Exogenous manipulation of targets (e.g. by a firm) is not a necessary condition for reference-dependence with goal-setting.

I also included in the experiment questions after the task about strategies they may have used. I also asked whether they enjoyed the task and whether they felt stress during the task. These questions allow respondents to state their internal thought processes about their observed performance each round as well as a proxy of experienced utility and disutility. This experiment does not explicitly test for present bias, but assuming any degree of bias, the randomization rules out imbalances in underlying present bias across treatment conditions and, therefore, isolates differential responses to the compensation incentives.

8 Experiment Results

I now present the results from my experiment from each round in Figure 6. Here, I primarily focus on the comparison between the bonus payment at 2,000 points and the piece rate condition, which has shown comparable mean performance in prior studies (DellaVigna and Pope, 2017).[33] Panel A of Figure 6 shows the distribution of performance in the first round comparing the piece rate treatment and the bonus condition at 2,000 points. Panel B shows the same for the subsequent

[33]Figure A9 shows a round-by-round comparison of the bonus condition at 2,400 points and the piece rate and exhibits remarkably similar dynamics.

rounds.

Several notable patterns emerge. First, there is substantial bunching in the distribution of performance. Importantly, the distributions in the bonus condition exhibit heaping to the right of the piece rate condition in every round. Panel C presents the differences in the densities of round 1 performance between the two conditions and shows that the bonus conditions are heavily concentrated above the 500- and 600-point thresholds. Next, despite having the same payoff at 2,000 points, the distribution of performance is consistently narrower in the bonus condition than in the piece rate condition, not only in the distribution of end performance but for the first three rounds. Because the only difference between the two groups is the compensation condition, the bonus condition itself is *causally* inducing participants to adopt (or increase their use of) round-specific references.[34]

To summarize the difference between the bonus condition and the piece rate condition more concretely, Figure 7 shows the density of average per-round performance in 5-point bins for the two conditions for the first three rounds (Panel A) and for all four rounds (Panel B). 500 points is the average each round that those in the bonus condition would have to perform in order to achieve the bonus at the end of the fourth round. Each side of the 500-point cutoff is approximated with a simple quadratic function of the density, and the solid lines denote the excess mass accruing to the right side of the cutoff. From this summary measure, it is visually clear that the missing mass between 400 and 475 points is eclipsed by the excess mass in the bonus condition between 500 and 550 points, after which the smoothed distributions are nearly identical up to 600 points.

Formally, I calculate excess mass in rounds 1-3 by comparing the estimated densities using these quadratic splines, i.e. measuring the ratio of the estimated densities in the 500-600 point range. This results in an excess mass of 9.4% (with a bootstrapped standard error of 0.504%) in the bonus condition. This is comparable to the 9.2% - 11.7% excess mass among loss-averse electronic tax filers around a zero final balance in Engström et al. (2015). Thus, workers in the 2,000-point condition are engaged in considerable bunching around this threshold every round as a result of the non-linear payment mechanism even though both conditions have the same ex ante expected payout at the same expected mean performance. In other words, even in the first three rounds when the exact performance in each round is not payoff-relevant, participants are exhibiting loss aversion in an effort to surpass *at least* the average required performance each round to get their bonus.

Another important observation is that the left tail in the density of round-specific performance increases in mass substantially in the final round and does so only in the bonus condition (and the same is true in Figure A9 for the bonus at 2,400 points). The result is a steep increase in the variance of effort during the final round in the bonus condition, whereas the variance does not

[34]In the piece rate condition, participants may be anchoring to round numbers. That there is more heaping in the bonus conditions is important because the piece rate represents a counterfactual that takes into account any tendency to bunch at round numbers.

exhibit the same behavior in the piece rate scheme. These patterns are clear in Figure A10. The variance in the piece rate condition is much higher during the first three rounds, and the changes in the variance across rounds are nearly perfectly parallel until the final round. During the final round, the standard deviation of performance rises by nearly 35% in the bonus condition.

The panels of Figure 6 demonstrate other interesting results. There appear to be two heaping points at 500 points and 600 points and this heaping is more pronounced in the bonus conditions.[35] In Figure A9, the heaping points are higher when the bonus is set at 2,400 compared to 2,000 points (closer to 700), meaning that round-specific sub-goals are responsive to the location of the end goal. When respondents were asked an open-ended question about their strategy, many in the 2,000-point treatment responded they targeted 600 points for each round as a buffer against fatigue or surprises in later rounds in order to avoid missing the 2,000-point threshold. Similarly, many of those in the 2,400-point bonus treatment stated that 700 was their round-specific target as a hedge against risk.

Panel D of Figure 6 shows the total performance across the three experimental conditions. These densities affirm the predictions discussed above in the sales data. The variance of total performance is lower in the bonus conditions than in the piece rate condition. Consistent with the Kőszegi and Rabin (2006) model, performance is higher with the bonus at 2,400 than the bonus at 2,000. Interestingly, despite cumulative performance of 2,000 points not being relevant at all to payoffs in the 2,400-point bonus condition, there is still substantial distributional heaping at 2,000 total points. Thus, non-linear payments create or make salient personal targets, even if those targets are not immediately relevant for final payoffs.

These results have significant implications for firm costs. In the experiment, the average bonus payouts for the piece rate condition were \$1.15 per worker compared to \$0.80 for the bonus condition at 2,000 points, representing a statistically significant reduction in per-person costs of 31%. Meanwhile, the average total output for the 2,000-point bonus condition was 1.58% lower and not statistically different from the piece rate condition. In the 2,400 point bonus, total payouts were only \$0.62 per worker, while output was slightly and marginally significantly higher than the piece rate. Thus, consistent with the Koch and Nafziger (2020) model of increased performance and goal attainment through short-run reference dependence, the core reason for these differences is attributable to the bonus condition leading to loss aversion in each round, which increases the likelihood of attaining a minimum performance threshold in the "long run" at a lower cost. This approach to establishing expectations via non-linear payments is advantageous to the firm.

One strength of this online experiment is the ability to directly elicit measures of enjoyment and disutility (proxied by feeling stress during the task). This would help separate actual reference

[35]Payoffs above 2,000 points are also *higher* in the piece rate condition, so more bunching at 600 (above what would be necessary to achieve 2,000 total points if that performance continued) in early rounds in the bonus condition is notable.

dependence from planning heuristics or other non-utility-based models of behavior. The patterns of responses across the distribution of performance are informative, particularly because reference dependence with loss aversion implies lower total utility just below the reference point. Figure A11 reveals that Round 1 performance relative to benchmarks of 500 and 600 points much more strongly predicts enjoyment of the task and stress under the bonus condition (at 2,000) relative to the piece rate condition (Panels A and B). This relationship holds even when nonparametrically controlling for total performance at the end of the task (Panels C and D). Conditional on cumulative performance, Round 1 performance is not payoff relevant in the bonus condition, yet there is still a substantial gap in the distributions of enjoyment and stress reported at the *end* of the task at 500 and 600 points. This lack of enjoyment and increase in stress below the reference point is consistent with loss aversion for these targets despite their lack of payoff relevance.[36] A planning heuristic or other non-utility-based models of this behavior would not generate this pattern of enjoyment or stress around these cutoffs when conditioning on performance.

Finally, after the end of the task period, I asked each participant an open-ended question: "Did you have any particular strategy when performing the task across these rounds?" Most participants shared information about their hand placement or other physical movements. However, 35 respondents explicitly stated unprompted that they had an internal target of 500 per round, 29 of which were presented with the bonus conditions. An additional 14 bonus participants pinpointed 600 as their target compared to only two in the piece rate, meaning that bonus condition participants were nearly five times more likely to articulate this type of targeted goal-setting as their primary, salient focus across rounds.[37]

Taken together, these results make clear two key empirical observations from a controlled experimental setting. First, workers do exhibit reference-dependent preferences over short-run performance as a means of holding themselves accountable for achieving a certain level of performance. Both revealed effort and explicit declarations of their internal thought processes confirm this result. Second, firms (i.e. the experimenter) can *causally induce* such behavior by using nonlinear payment schemes to make long-run and short-run reference points salient and meaningful. Inducing such short-term goal-setting via these payment schemes generates substantial performance improvements relative to the monetary costs by inducing workers to impose psychological costs upon themselves in the form of loss aversion.

9 Discussion and Conclusion

Using novel data from a door-to-door sales company and an online experiment, this paper provides evidence of reference-dependent preferences in daily labor supply. Door-to-door sales

[36]Imperfect ability to predict final performance at the end of round 1 may be a reason for the enjoyment gap persisting despite narrowly missing 500.

[37]Some examples of responses include: "I ... was trying to get over 600 in the first rounds since I knew my fingers would be tired by the last round." "I made it a goal to get to 500 on each." "[I went] as fast as I possibly could ... to go over the 500 mark per round for the first three rounds."

workers exhibit loss aversion around expectations-based references in their labor supply choices. I find that the extensive margin choice (when to stop working) is the margin at which reference dependence is most operative. I show that by making particular points in the final performance distribution salient and consequential, the firm's bonus schedule for end-of-season sales facilitates this behavior: workers impose upon themselves daily goals and exhibit loss aversion around these targets in the service of attaining the bonus. Upon surpassing the relevant bonus threshold, they drastically reduce their labor supply even when they are still paid a significant piece rate.

My online real-effort task experiment confirms the causal interpretation of the sales data and shows that by simply having a bonus payment that makes a particular long-run target salient, workers respond by imposing upon themselves short-run targets that serve as reference points. They do this as a commitment device in order to achieve their (and the firm's) larger goals. That firms can induce this type of short-run loss averse has not been demonstrated before in the literature.

This result provides new information about *why* workers might exhibit reference dependence in their labor supply choices—as a rational response to expected present bias. Loss aversion around daily goals acts as a commitment device, keeping these workers engaged in work when they might otherwise shirk today assuming they can catch up tomorrow. Narrow bracketing of goals, therefore, frequently induces effort by keeping workers in a loss domain at the start of each day, leveraging one behavioral bias (loss aversion) to overcome another (dynamic inconsistency) (Heath et al., 1999; Koch and Nafziger, 2016; Hsiaw, 2018; Koch and Nafziger, 2020), and workers self-impose these preferences. These results provide a very clear dynamic mechanism for why non-linear compensation schemes are effective in raising worker output (e.g. Freeman and Gelber (2010); Imas et al. (2017); Kuhn and Yu (2021)): workers adjust their short-run labor supply choices to include loss aversion, which combats present bias. This also explains why people might choose dominated daily contracts (Kaur et al., 2015). Unlike in prior settings, this daily loss aversion does not require costly real-time observation of daily performance or loss-framed contracts. These results also offer an explanation for the daily income targeting observed in the taxi literature (e.g. Camerer et al. (1997); Thakral and Tô (2021)) and in other settings (e.g. Dupas et al. (2020)). Given the unpleasantness of door-to-door sales, self-control problems are likely universal in the occupation. These results also reveal a key mechanism behind the effects of non-linear incentives in the workplace: the establishment of expectations, both for the short and long run (Kőszegi and Rabin, 2006). The online experiment confirms that compensation schemes used by firms can "make" or reinforce reference dependence.

These results are broadly applicable across many occupational contexts. The types of incentives in my experiment—non-linear bonuses and piece rates—are common features of a wide variety of occupations. These incentives are widely used in sales, which represents a large global market. The behaviors of door-to-door sellers, therefore, can easily be generalized to other sales and marketing occupations. Other industries and labor markets use these types of incentives. Piece rates are common in many occupations where outcomes can be objectively measured, from fruit picking (Graff-Zivin et al., 2019) to investment commissions for financial managers. The use of formal and informal bonuses at performance targets is ubiquitous, from the highest-paid CEOs to children selling coupon books to raise money for their sports or performing arts programs. That these are widely used across occupations and contexts indicates that a diverse set of actors acknowledge the motivational power of these incentives.

My results have important implications for how workers optimize their labor supply. Because workers are more motivated by additional income in the loss domain and less motivated in the gain domain, the effectiveness of a wage increase depends on the worker's reference point. My results suggest that the firm, rather than trying to motivate *around* a reference point, can influence the *positioning* of the reference point itself. These results have significant implications for firm efficiency as the online experiment shows that the bonus conditions produced significant cost savings relative to worker output. From the firm's perspective, inducing narrow goals in this way is a low-cost *psychological* alternative to high-cost *monetary* incentives. This may partially explain the use of these non-linear incentives across many workplaces.

References

Agarwal, Sumit, Mi Diao, Jessica Pan, and Tien Foo Sing. 2015. "Are Singaporean Cabdrivers Target Earners?" *Available at SSRN 2338476*.

Allen, Eric J, Patricia M Dechow, Devin G Pope, and George Wu. 2017. "Reference-Dependent Preferences: Evidence from Marathon Runners." *Management Science*, 63(6): 1657–1672.

Andersen, Steffen, Alec Brandon, Uri Gneezy, and John A List. 2014. "Toward an Understanding of Reference-Dependent Labor Supply: Theory and Evidence from a Field Experiment."Technical report, NBER.

Ariely, Dan, and Klaus Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological science*, 13(3): 219–224.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *The Quarterly Journal of Economics*, 121(2): 635–672.

Bellemare, Charles, and Bruce Shearer. 2010. "Sorting, Incentives and Risk Preferences: Evidence from a Field Experiment." *Economics Letters*, 108(3): 345–348.

Bénabou, Roland, and Jean Tirole. 2004. "Willpower and Personal Rules." *Journal of Political Economy*, 112(4): 848–886.

Bettinger, Eric P, Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block FAFSA Experiment." *The Quarterly Journal of Economics*, 127(3): 1205–1242.

Brown, Alexander L., Zhikang Eric Chua, and Colin F. Camerer. 2009. "Learning and Visceral Temptation in Dynamic Saving Experiments*." *The Quarterly Journal of Economics*, 124(1): 197–231, URL: https://doi.org/10.1162/qjec.2009.124.1.197, DOI: http://dx.doi.org/10.1162/qjec.2009.124.1.197.

Cadena, Brian C, and Benjamin J Keys. 2015. "Human Capital and the Lifetime Costs of Impatience." *American Economic Journal: Economic Policy*, 7(3): 126–53.

Cai, Xiqian, Wei Jiang, Hong Song, and Huihua Xie. 2022. "Pay for Performance Schemes and Manufacturing Worker Productivity: Evidence from a Kinked Design in China." *Journal of Development Economics*, 156, p. 102840.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *The Quarterly Journal of Economics*, 112(2): 407–441.

Cawley, John, and Christopher J Ruhm. 2011. "The Economics of Risky Health Behaviors." In *Handbook of Health Economics*. 2: Elsevier, 95–199.

Chou, Yuan K. 2002. "Testing Alternative Models of Labour Supply: Evidence from Taxi Drivers in Singapore." *The Singapore Economic Review*, 47(01): 17–47.

Correia, Sergio. 2016. "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator."Technical report, Working Paper.

Crawford, Vincent P., and Juanjuan Meng. 2011. "New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income." *American Economic Review*, 101(5): 1912–1932, DOI: http://dx.doi.org/10.1257/aer.101.5.1912.

Dellavigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315–372, DOI: http://dx.doi.org/10.1257/jel.47.2.315.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review*, 96(3): 694–719.

DellaVigna, Stefano, and Devin Pope. 2017. "What Motivates Effort? Evidence and Expert Forecasts." *The Review of Economic Studies*, 85(2): 1029–1069.

Dupas, Pascaline, Jonathan Robinson, and Santiago Saavedra. 2020. "The Daily Grind: Cash Needs and Labor Supply." *Journal of Economic*

Behavior & Organization, 177 399–414.
Engström, Per, Katarina Nordblom, Henry Ohlsson, and Annika Persson. 2015. “Tax Compliance and Loss Aversion.” *American Economic Journal: Economic Policy*, 7(4): 132–64.
Farber, Henry S. 2005. “Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers.” *Journal of Political Economy*, 113(1): 46–82.
Farber, Henry S. 2008. “Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers.” *American Economic Review*, 98(3): 1069–1082, DOI: http://dx.doi.org/10.1257/aer.98.3.1069.
Farber, Henry S. 2015. “Why You Can’t Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers.” *The Quarterly Journal of Economics*, 130(4): 1975–2026, DOI: http://dx.doi.org/10.1093/qje/qjv026.
Fehr, Ernst, and Lorenz Goette. 2007. “Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment.” *American Economic Review*, 97(1): 298–317.
Fouarge, Didier, Ben Kriechel, and Thomas Dohmen. 2014. “Occupational Sorting of School Graduates: The Role of Economic Preferences.” *Journal of Economic Behavior & Organization*, 106 335–351.
Freeman, Richard B, and Alexander M Gelber. 2010. “Prize Structure and Information in Tournaments: Experimental Evidence.” *American Economic Journal: Applied Economics*, 2(1): 149–64.
Freeman, Richard B, Wei Huang, and Teng Li. 2019. “Non-linear Incentives and Worker Productivity and Earnings: Evidence from a Quasi-experiment.” *NBER Working Paper*(25507): , URL: http://www.nber.org/papers/w25507.
Gächter, Simon, Eric J Johnson, and Andreas Herrmann. 2007. “Individual-Level Loss Aversion in Riskless and Risky Choices.”
Giné, Xavier, Monica Martinez-Bravo, and Marian Vidal-Fernández. 2016. “Are Labor Supply Decisions Consistent with Neoclassical Preferences? Evidence from Indian Boat Owners.” URL: https://elibrary.worldbank.org/doi/abs/10.1596/1813-9450-7820, DOI: http://dx.doi.org/10.1596/1813-9450-7820.
Goette, Lorenz, David Huffman, and Ernst Fehr. 2004. “Loss Aversion and Labor Supply.” *Journal of the European Economic Association*, 2(2-3): 216–228.
Goldberg, Jessica. 2016. “Kwacha Gonna Do? Experimental Evidence About Labor Supply in Rural Malawi.” *American Economic Journal: Applied Economics*, 8(1): 129–49.
Graff-Zivin, Joshua S, Lisa B Kahn, and Matthew J Neidell. 2019. “Incentivizing Learning-By-Doing: The Role of Compensation Schemes.”Technical report, NBER.
He, Shu, Liangfei Qiu, and Xusen Cheng. 2018. “Wage Elasticity of Labor Supply in Real-Time Ridesharing Markets: An Empirical Analysis.” *University of Connecticut School of Business Research Paper*(18-21): .
Heath, Chip, Richard P Larrick, and George Wu. 1999. “Goals as Reference Points.” *Cognitive Psychology*, 38(1): 79–109.
Hsiaw, Alice. 2018. “Goal Bracketing and Self-Control.” *Games and Economic Behavior*, 111, DOI: http://dx.doi.org/10.1016/j.geb.2018.06.005.
Ichniowski, Casey, and Kathryn Shaw. 2003. “Beyond Incentive Pay: Insiders’ Estimates of the Value of Complementary Human Resource Management Practices.” *Journal of Economic Perspectives*, 17(1): 155–180.
Imas, Alex, Sally Sadoff, and Anya Samek. 2017. “Do People Anticipate Loss Aversion?” *Management Science*, 63(5): 1271–1284.
Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica*, 47(2): 263–291.
Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2015. “Self-Control at Work.” *Journal of Political Economy*, 123(6): 1227–1277.
Köbberling, Veronika, and Peter P Wakker. 2005. “An Index of Loss Aversion.” *Journal of Economic Theory*, 122(1): 119–131.
Koch, Alexander K, and Julia Nafziger. 2016. “Goals and Bracketing Under Mental Accounting.” *Journal of Economic Theory*, 162 305–351.
Koch, Alexander K, and Julia Nafziger. 2020. “Motivational Goal Bracketing: An Experiment.” *Journal of Economic Theory*, 185 104–149.
Kőszegi, Botond, and Matthew Rabin. 2006. “A Model of Reference-Dependent Preferences.” *The Quarterly Journal of Economics*, 121(4): 1133–1165.
Kőszegi, Botond, and Matthew Rabin. 2007. “Reference-Dependent Risk Attitudes.” *American Economic Review*, 97(4): 1047–1073.
Kuhn, Peter J, and Lizi Yu. 2021. “Kinks as Goals: Accelerating Commissions and the Performance of Sales Teams.”
Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting.” *The Quarterly Journal of Economics*, 112(2): 443–478.
Markle, Alex, George Wu, Rebecca White, and Aaron Sackett. 2018. “Goals As Reference Points in Marathon Running: A Novel Test of Reference Dependence.” *Journal of Risk and Uncertainty*, 56(1): 19–50.
Martin, Vincent. 2017. “When to Quit: Narrow Bracketing and Reference Dependence in Taxi Drivers.” *Journal of Economic Behavior and Organization*, 144 166–187, URL: http://dx.doi.org/10.1016/j.jebo.2017.09.024, DOI: http://dx.doi.org/10.1016/j.jebo.2017.09.024.
McClure, Samuel M, David I Laibson, George Loewenstein, and Jonathan D Cohen. 2004. “Separate neural systems value immediate and delayed monetary rewards.” *Science*, 306(5695): 503–507.
Meier, Stephan, and Charles Sprenger. 2010. “Present-Biased Preferences and Credit Card Borrowing.” *American Economic Journal: Applied Economics*, 2(1): 193–210.
Menne, Matthew J, Imke Durre, Bryant Korzeniewski, Shelley McNeal, Kristy Thomas, Xungang Yin, Steven Anthony, Ron Ray, Russell S Vose, Byron E Gleason et al. 2012. “Global Historical Climatology Network-Daily (GHCN-Daily), Version 3.” *NOAA National Climatic Data Center*, 10, p. V5D21VHZ, DOI: http://dx.doi.org/10.7289/V5D21VHZ.
Mitchell, Olivia S. 1988. “Worker Knowledge of Pension Provisions.” *Journal of Labor Economics*, 6(1): 21–39.
Morgul, Ender Faruk, and Kaan Ozbay. 2015. “Revisiting Labor Supply of New York City Taxi Drivers: Empirical Evidence from Large-Scale Taxi Data.” In *Transportation Research Board 94th Annual Meeting*. 15.
Nguyen, Quang, and Pingsun Leung. 2013. “Revenue Targeting in Fisheries: The Case of Hawaii Longline Fishery.” *Environment and Development Economics*, 18(5): 559–575.
O’Donoghue, Ted, and Charles Sprenger. 2018. “Reference-Dependent Preferences.” In *Handbook of Behavioral Economics: Foundations and Applications 1*. eds. by B Douglas Bernheim, Stefano DellaVigna, and David Laibson: Elsevier, , Chap. 1 1–78.
Oettinger, Gerald S. 1999. “An Empirical Analysis of the Daily Labor Supply of Stadium Vendors.” *Journal of Political Economy*, 107(2): 360–392.
Park, R Jisung, Joshua Goodman, Michael Hurwitz, and Jonathan Smith. 2020. “Heat and Learning.” *American Economic Journal: Economic Policy*, 12(2): 306–39.
Patnaik, Arpita, Joanna Venator, Matthew Wiswall, and Basit Zafar. 2020. “The Role of Heterogeneous Risk Preferences, Discount Rates, and Earnings Expectations in College Major Choice.” *Journal of Econometrics*.
Schmidt, Marc-Antoine. 2018. “The Daily Labor Supply Response to Worker-Specific Earnings Shocks.”Technical report, Working Paper.
Shefrin, Hersh M., and Richard H. Thaler. 1992. “Mental Accounting, Saving, and Self-Control.” In *Choice Over Time*. eds. by George Lowen-

stein, and Jon Elster: New York: Russell Sage Foundation Press.
Sheldon, Michael. 2016. "Income Targeting and the Ridesharing Market." *Unpublished manuscript. Available at: https://static1. squarespace. com/static/56500157e4b0cb706005352d*, 56, p. 1457131797556.
Stafford, Tess M. 2015. "What Do Fishermen Tell Us that Taxi Drivers Do Not? An Empirical Investigation of Labor Supply." *Journal of Labor Economics*, 33(3): 683–710.
Thakral, Neil, and Linh T Tô. 2021. "Daily Labor Supply and Adaptive Reference Points." *American Economic Review*, 111(8): 2417–43.
Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.
Warner, John T, and Saul Pleeter. 2001. "The Personal Discount Rate: Evidence from Military Downsizing Programs." *American Economic Review*, 91(1): 33–53.

Figures

Figure 1: Contract Structure: Total Income by Sales ($500 Contract Value)

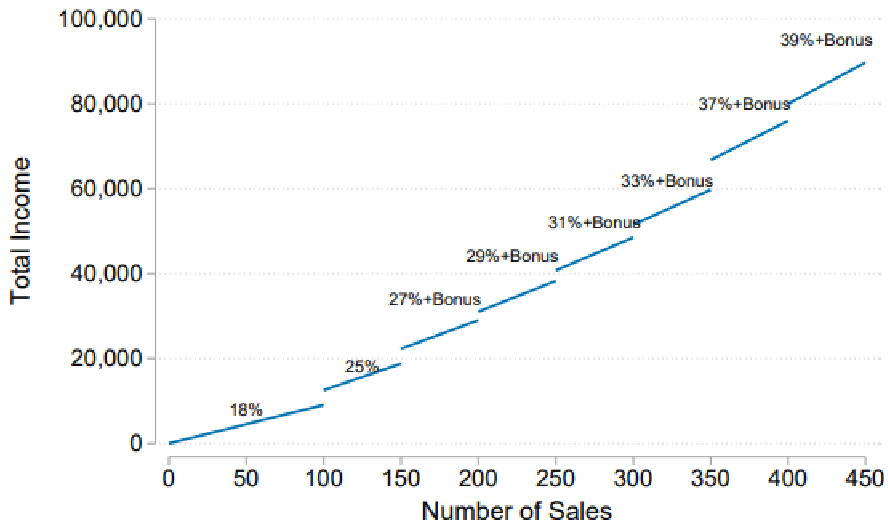

Source: Author's calculations of typical contracts from a pest control sales company.
Notes: Percentages indicate commissions as they are applied to each interval for all sales at the end of the season. At 150 sales, the "bonus" is that the company pays for the seller's rent for the summer in full (≈$2,000). At 250 sales, sellers qualify for the all-expenses-paid company vacation.

Figure 2: Labor Supply Around Expectations
Panel A: Probability of Stopping for the Day

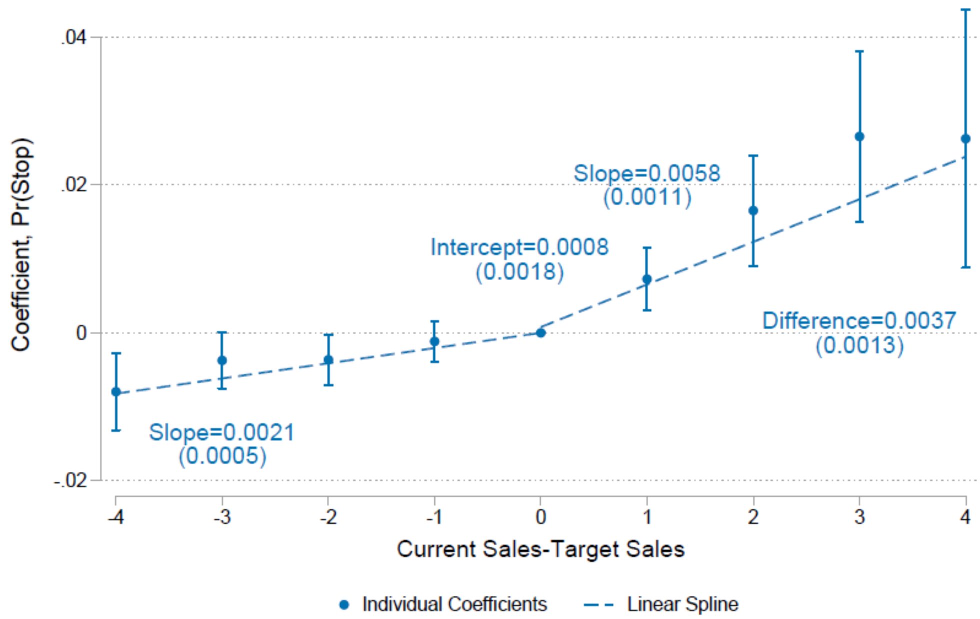

Panel B: Pitches Per Half Hour

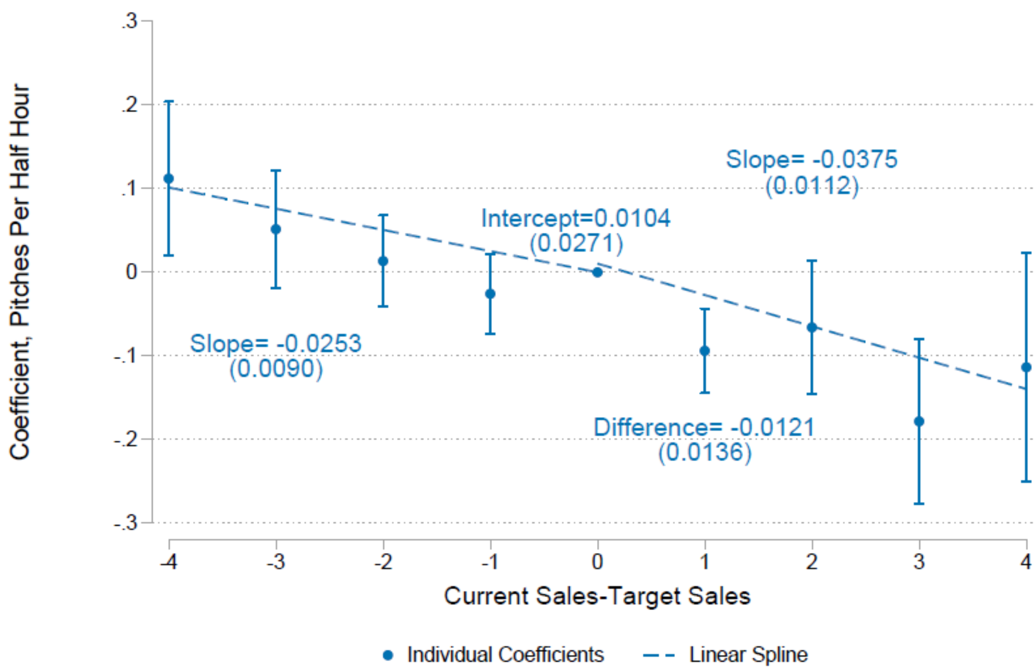

Source: Author's calculations of data from a pest control sales company.
Notes: Results are from estimates of Equations 1 and 2. Standard errors are clustered at the seller level. Individual coefficients are rounded to the nearest integer. The coefficient on +1 includes the interval (0,1.5).

Figure 3: Kernel Density of Total Sales by Week

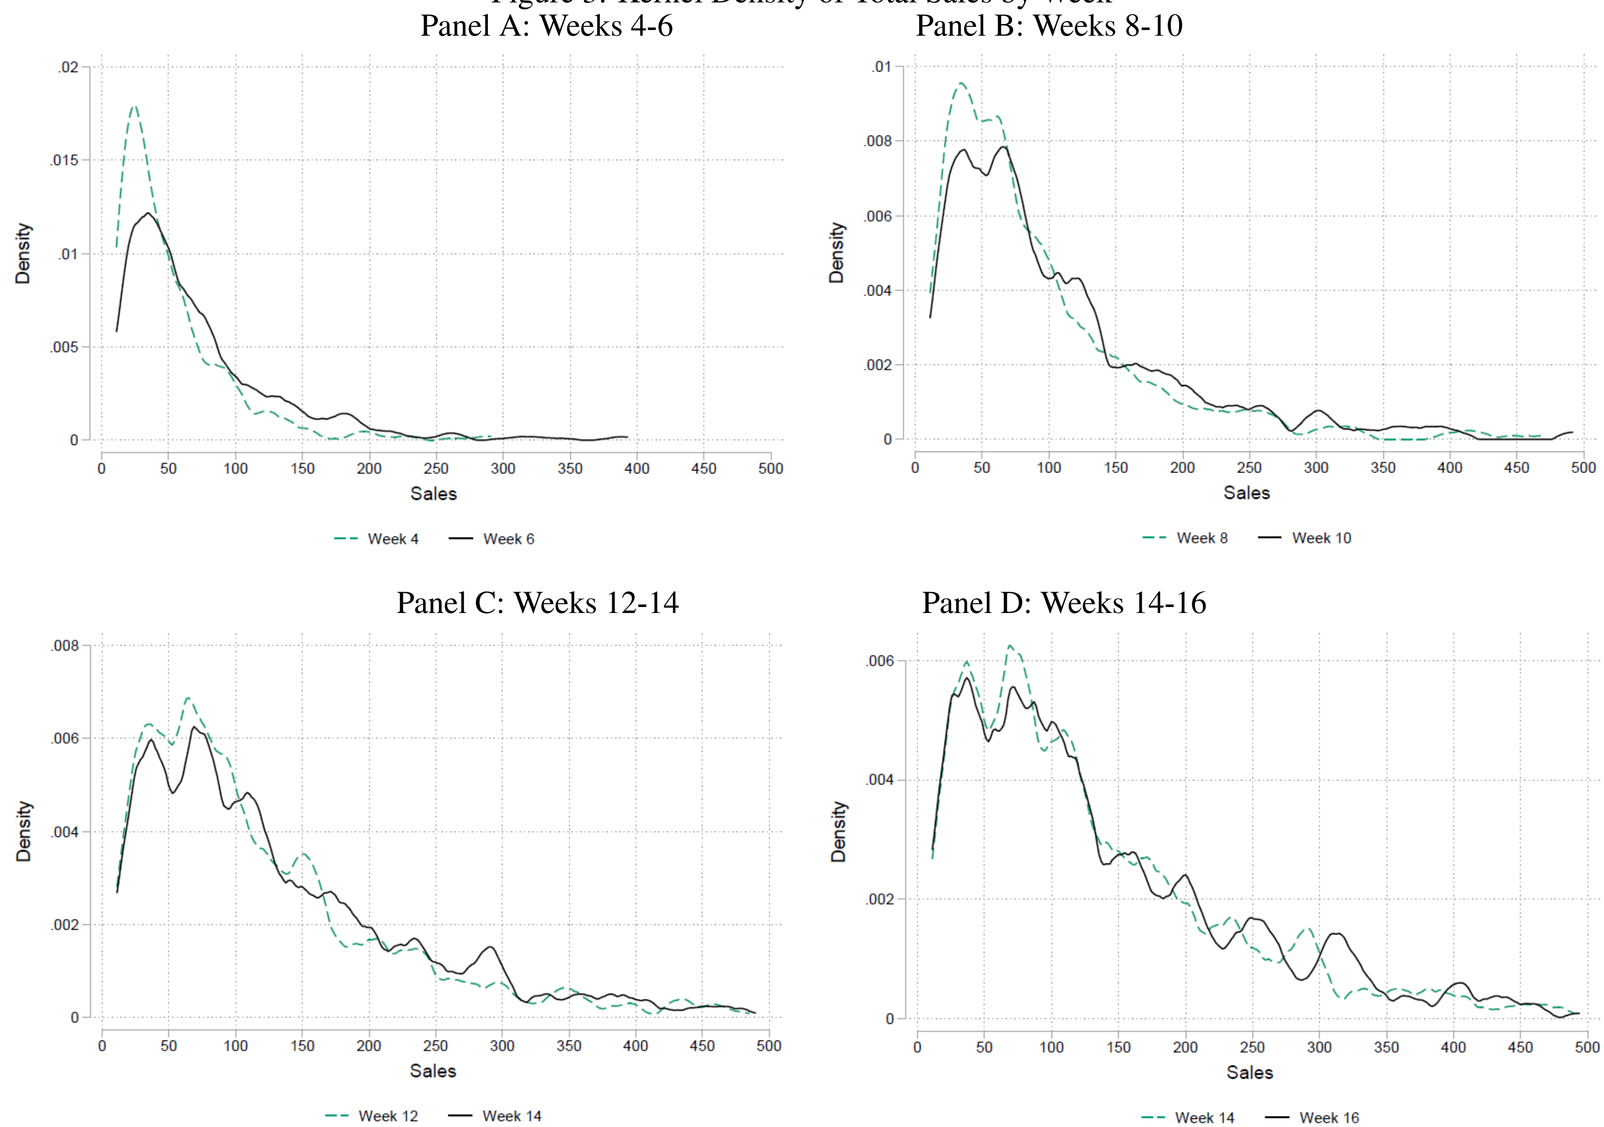

Source: Author's calculations of data from a pest control sales company.

Notes: Estimated using Epanechnikov kernel with bandwidth of 7 sales at the end of each estimated week among those with at least ten sales.

Figure 4: Kernel Density of Total Sales by Week
Workers with Total Sales of 175–225 at End of Season

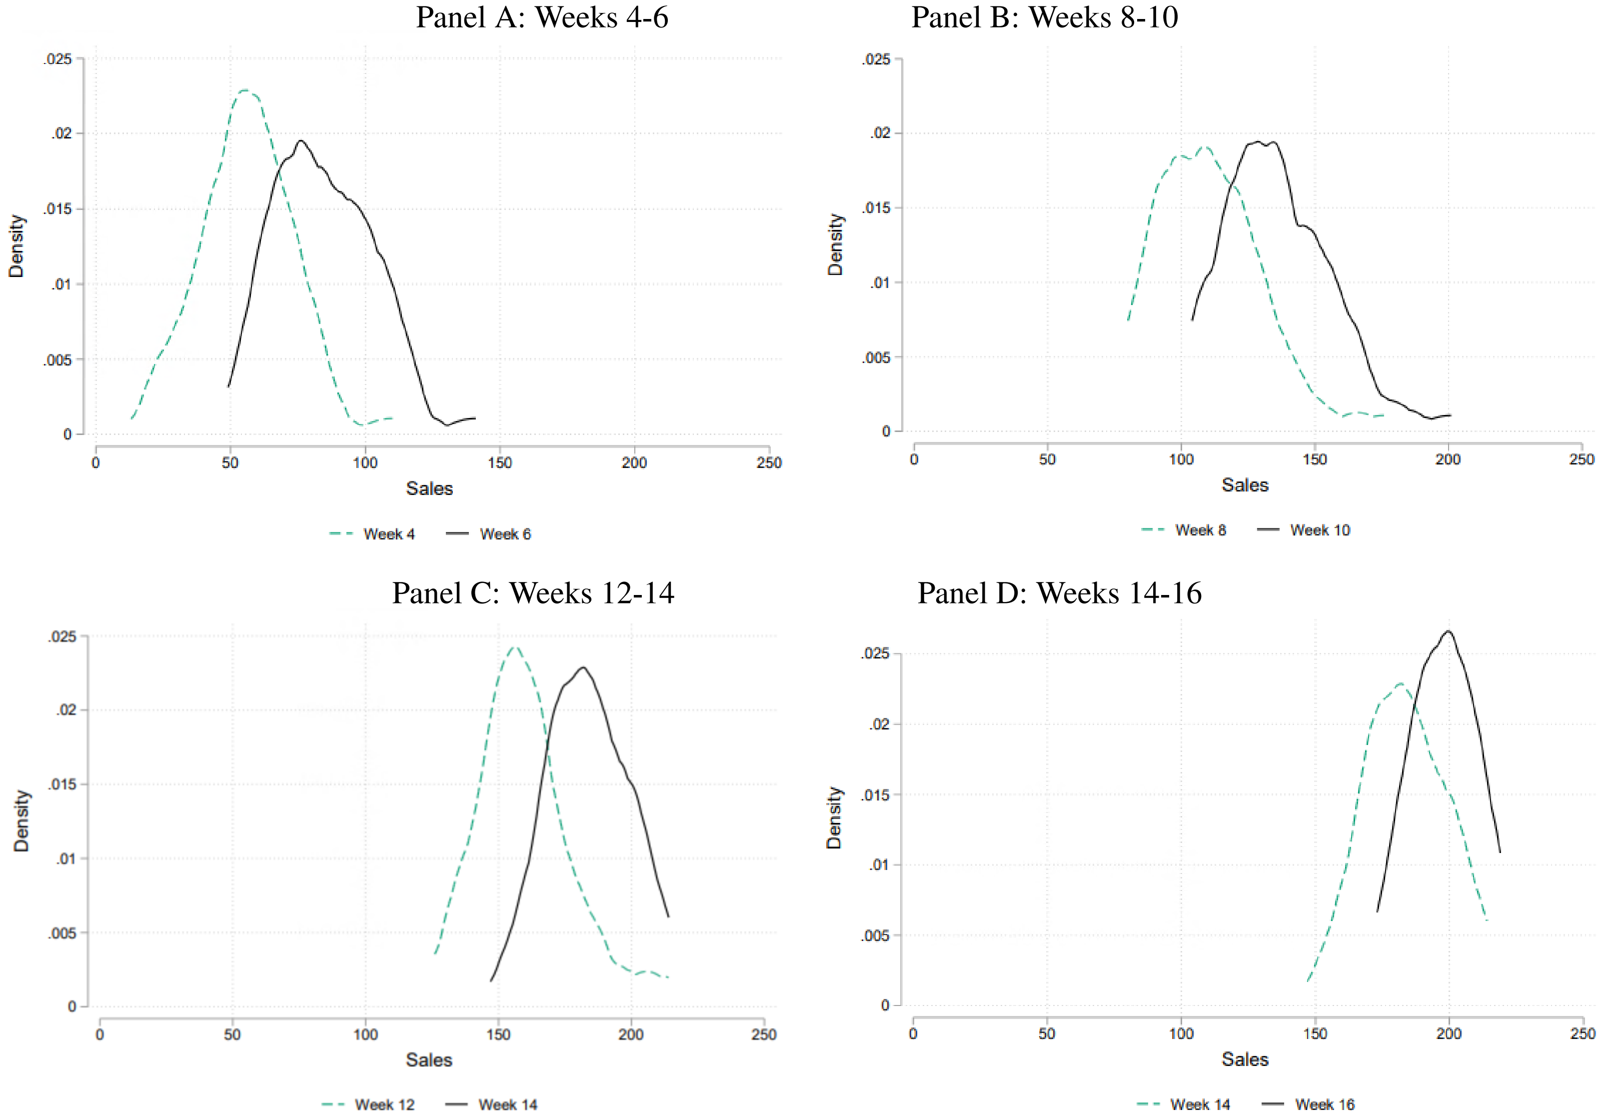

Source: Author's calculations of data from a pest control sales company.
Notes: Estimated using Epanechnikov kernel with bandwidth of 7 sales at the end of each estimated week among those with at least ten sales.

Figure 5: Predicted Labor Supply Over Current Sales Interval, By Final Season Sales Interval

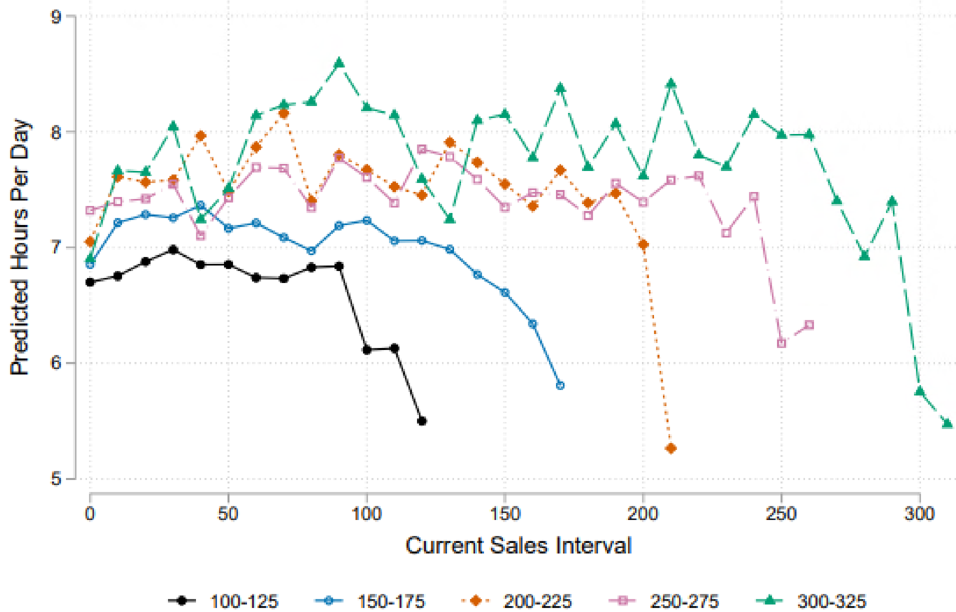

Source: Author's calculations of data from a pest control sales company.
Notes: Plot shows predicted hours from specification in Equation 3 for current sales interval (x-axis) separated by bins of total end-of-season sales.

Figure 6: Density Estimates of Experimental Performance

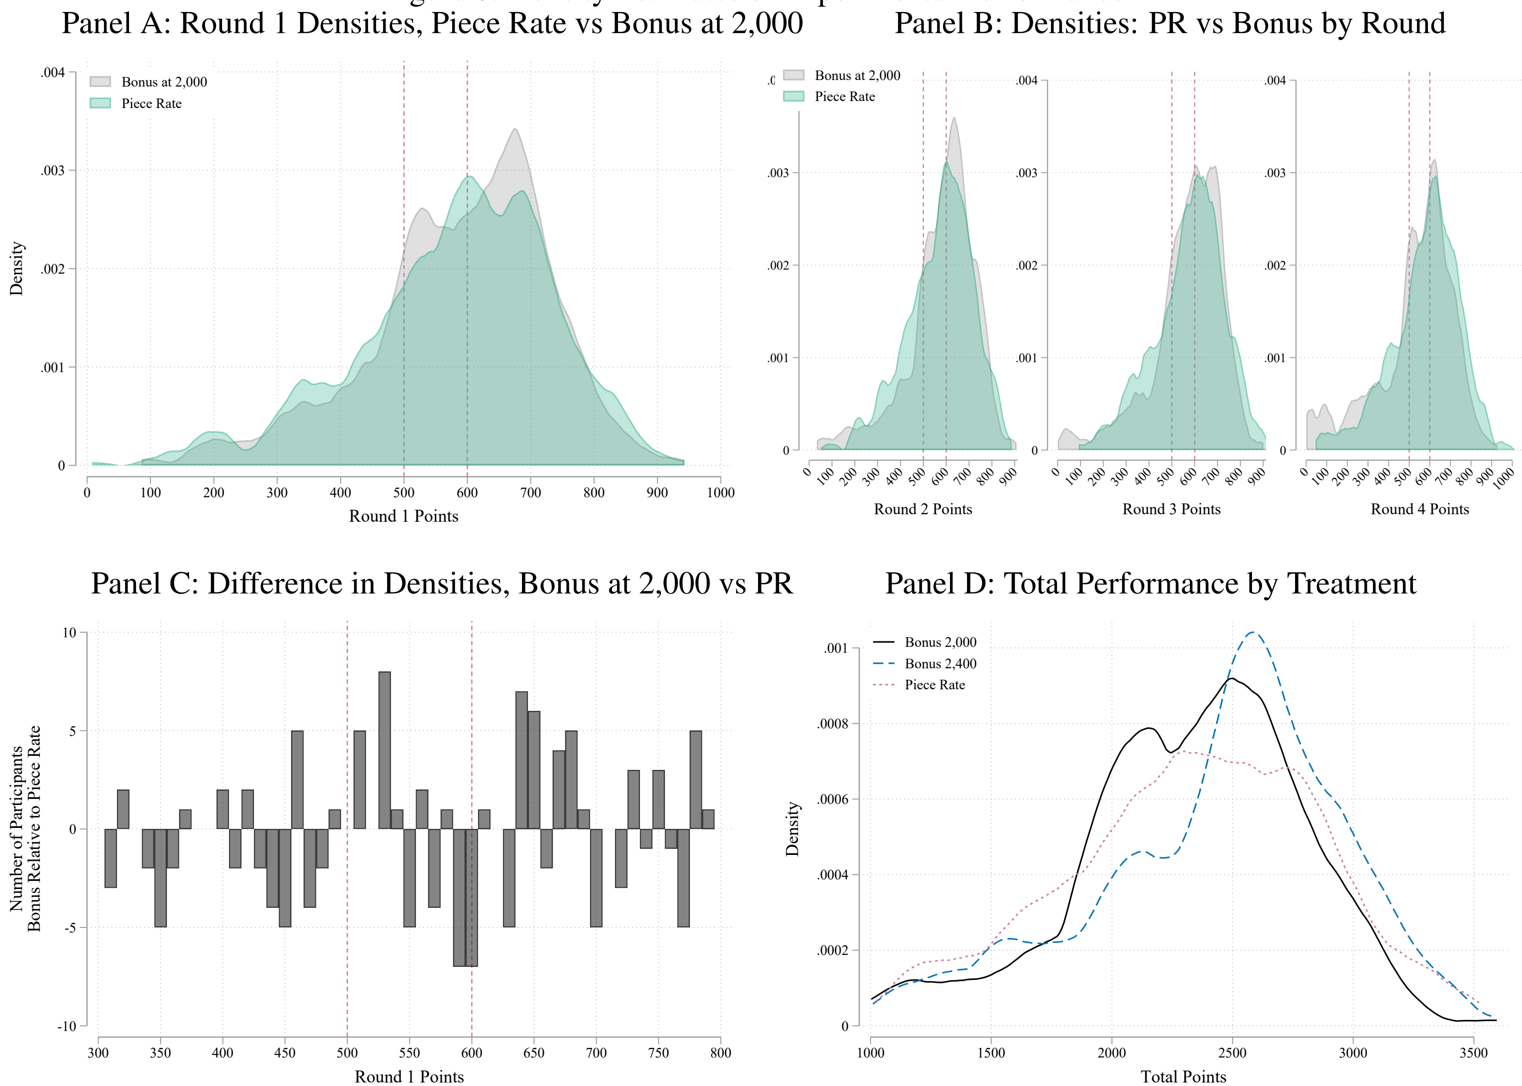

Source: Author's calculations of data from an online experiment.

Notes: Panel C breaks frequencies down into discrete 10-point categories and presents the differences between the distributions in Panel A. In Panel C, the X-axis values reflect the minimum of each 10-point interval.

Figure 7: Density Discontinuities at 500 Points Per Round

Panel A: Rounds 1-3

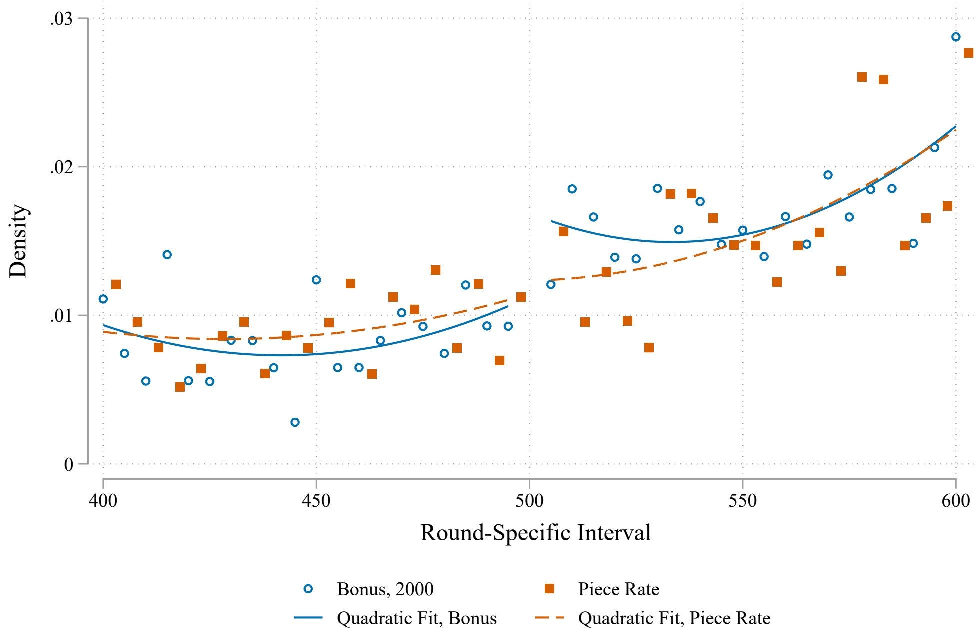

Panel B: All Rounds

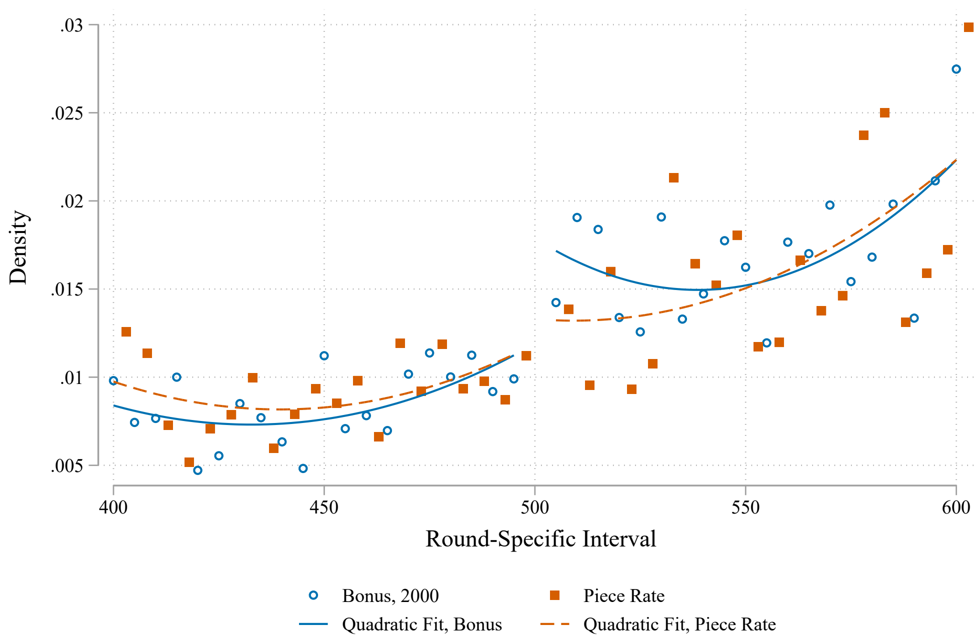

Source: Author's calculations of data from an online experiment.
Notes: The X-axis shows the average number of points achieved per round in the bonus condition and the piece rate conditions in 5-point intervals, while the Y-axis is the density. 500 points is the average each round that those in the bonus condition would have to perform in order to achieve the bonus at the end of the fourth round. Each side of the 500-point cutoff is approximated with a quadratic function of the density, and the higher solid blue lines denote the excess mass accruing to the right side of the cutoff. Using the smoothed estimates for the bonus versus the piece rate condition yields excess mass of 9.4% (bootstrapped standard error of 0.504%).

A Online Appendix: Figures and Tables

Figure A1: Distribution of Start and Stop Characteristics
Panel A: Time of Day at Start of Shift

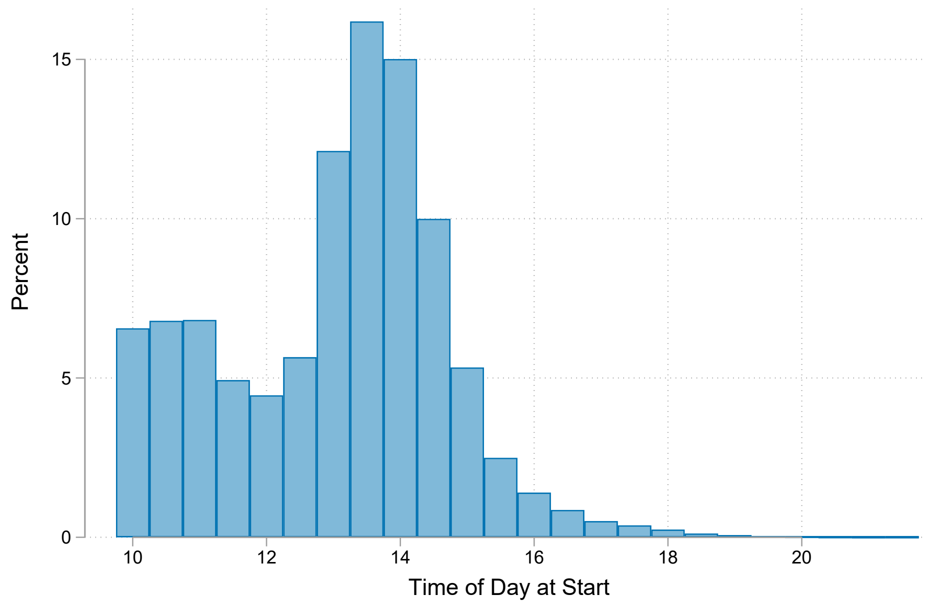

Panel B: Hour of Shift at Stop

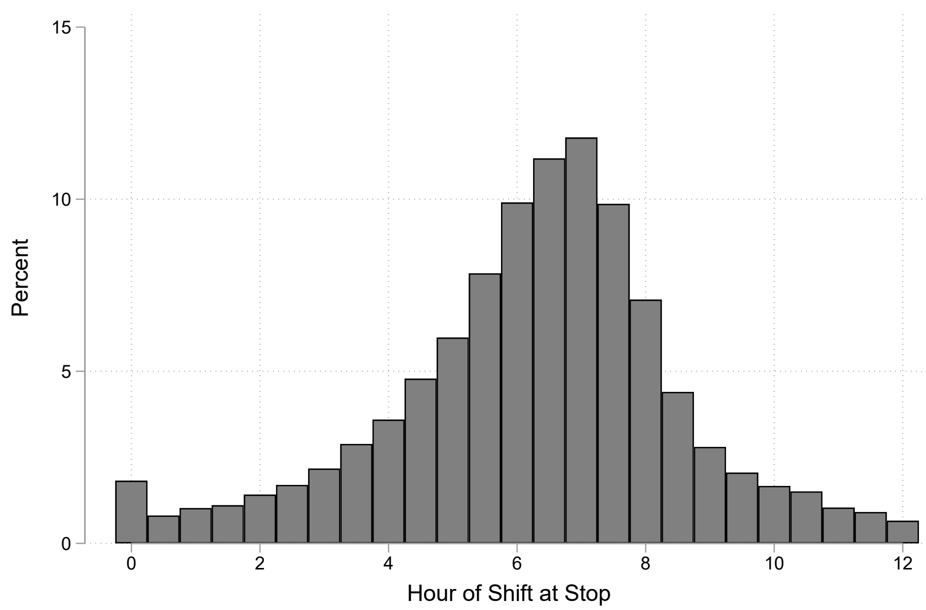

Source: Author's calculations of data from a pest control sales company.
Notes: Shifts begin during the half hour period when a seller first registers a knock or sale on each workday. Shifts end during the half hour they record their last sale or knock for the day.

Figure A2: Reference Dependence and Probability of Active Work

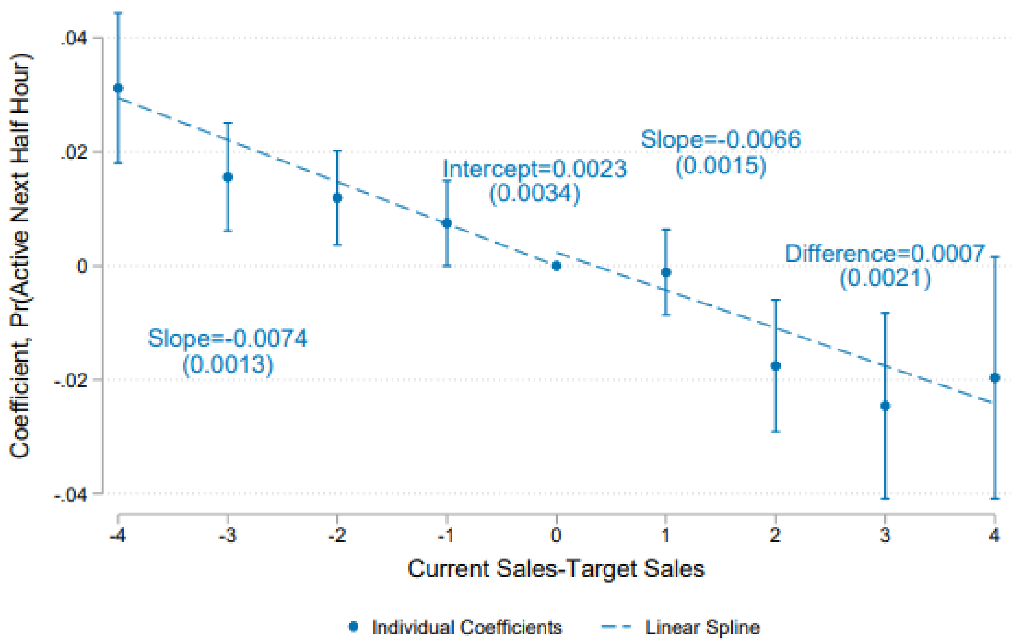

Source: Author's calculations of data from a pest control sales company.
Notes: Results are from estimates of Equations 1 and 2 for the probability of working during the next half hour of the day. Standard errors are clustered at the seller level. Individual coefficients are rounded to the nearest integer. The coefficient on +1 includes the interval (0,1.5). Model also includes an additional control for actively knocking during the current half-hour period. At a base active knocking share of 80% of all half-hour periods, an increase of 1 sale above or below expectations decreases the probability of actively knocking by approximately 0.74 percentage points, or approximately 0.093%. Sellers are not more likely to take breaks during their work as a function of their position relative to expectations.

Figure A3: Kernel Density of Total Sales at End of Season

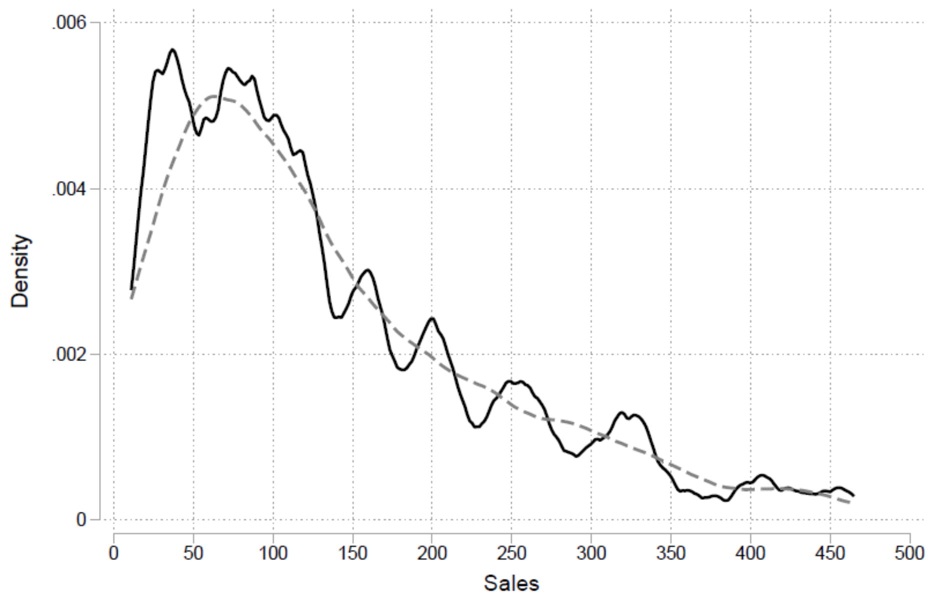

Source: Author's calculations of data from a pest control sales company.
Notes: Estimated using Epanechnikov kernel with bandwidth of 7 sales and 25 sales for sellers with at least ten sales and fewer than 500. The retroactive nature of the commission increases leads to a cash bonus upon hitting each 50-sale interval. At 150 sales, the company pays for the seller's rent for the summer in full (≈$2,000). At 250 sales, sellers qualify for the all-expenses-paid company vacation.

Figure A4: Kernel Density of Total Sales by Week
Workers with Total Sales of 125–175 at End of Season

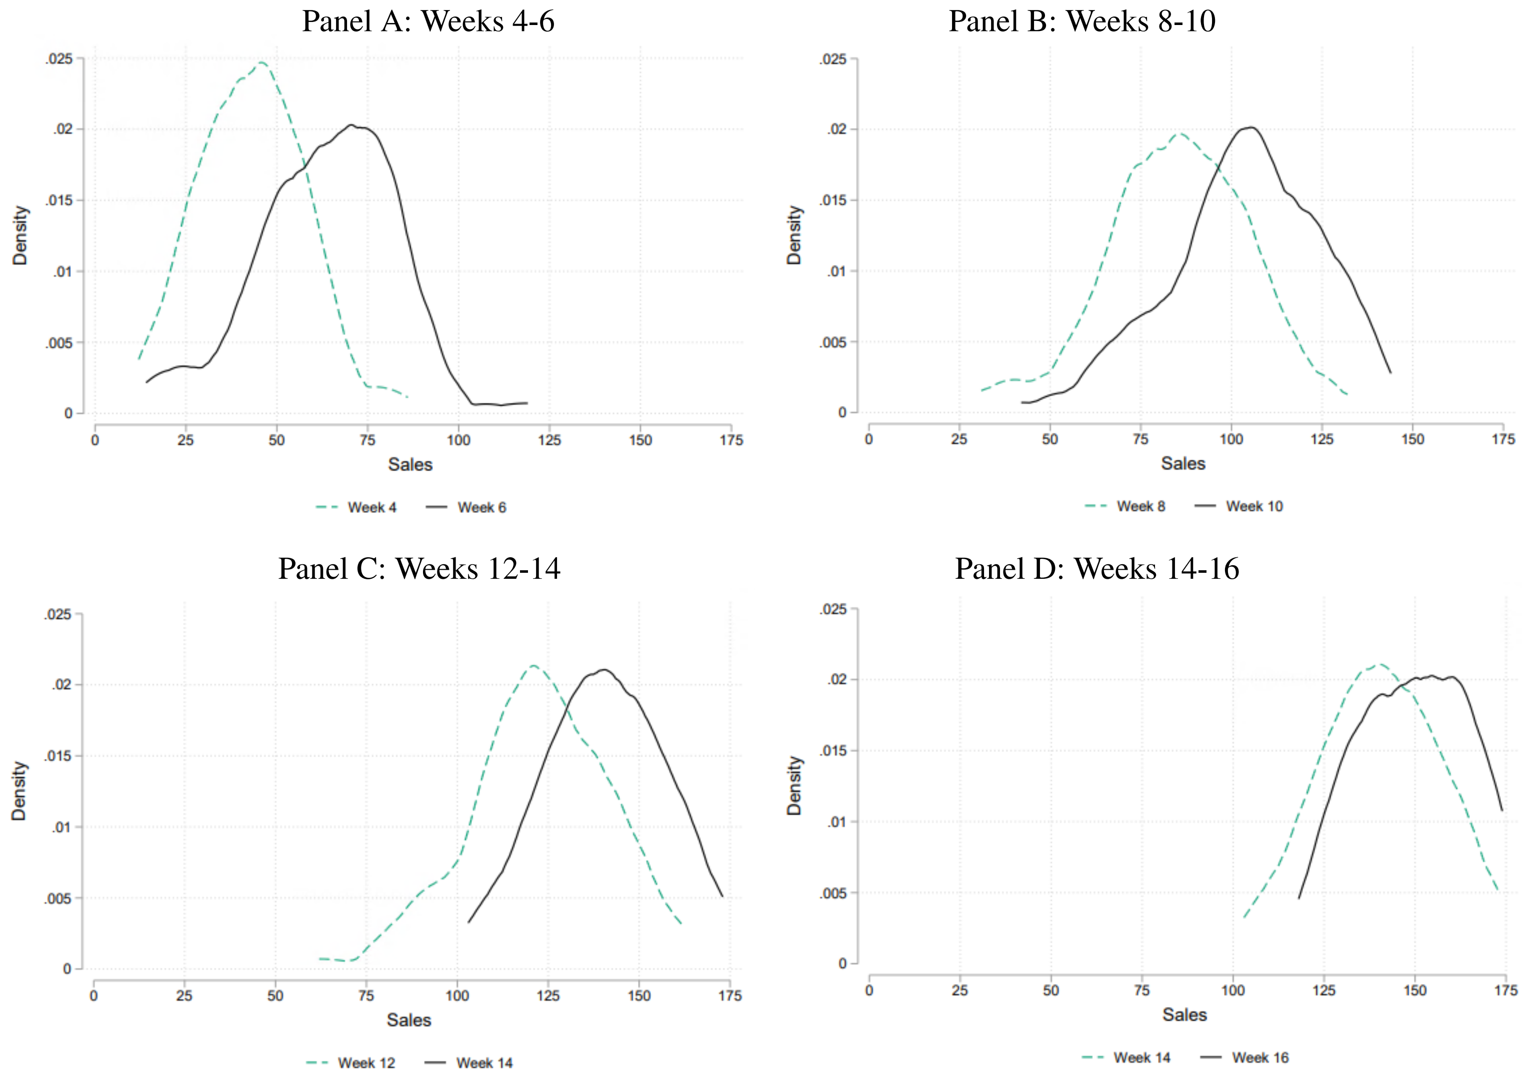

Source: Author's calculations of data from a pest control sales company.
Notes: Estimated using Epanechnikov kernel with bandwidth of 7 sales at the end of each estimated week among those with at least ten sales.

Figure A5: Standard Deviation of Residual Daily Hours Over Time and Progress Toward Total

Panel A: Hours Over Time, All Sellers

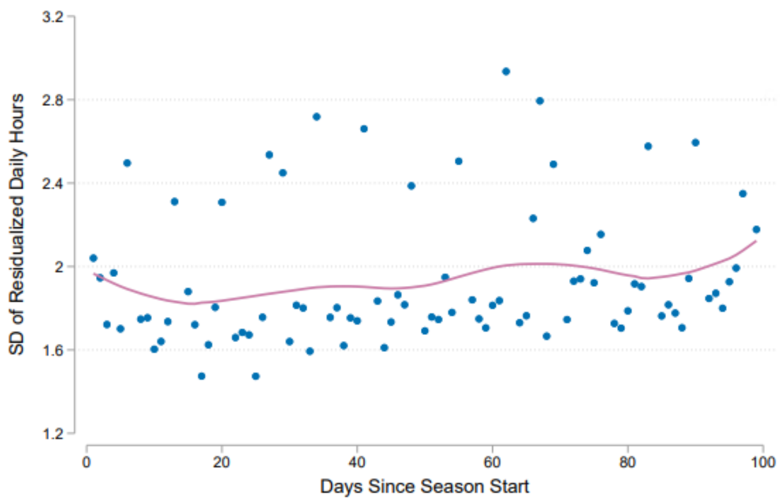

Panel B: Hours Over Progress Toward Total, All Sellers

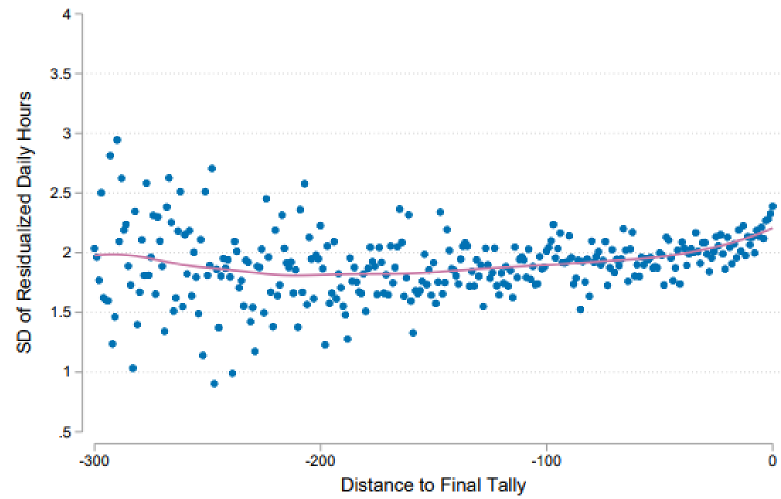

Panel C: Hours Over Time, by Total

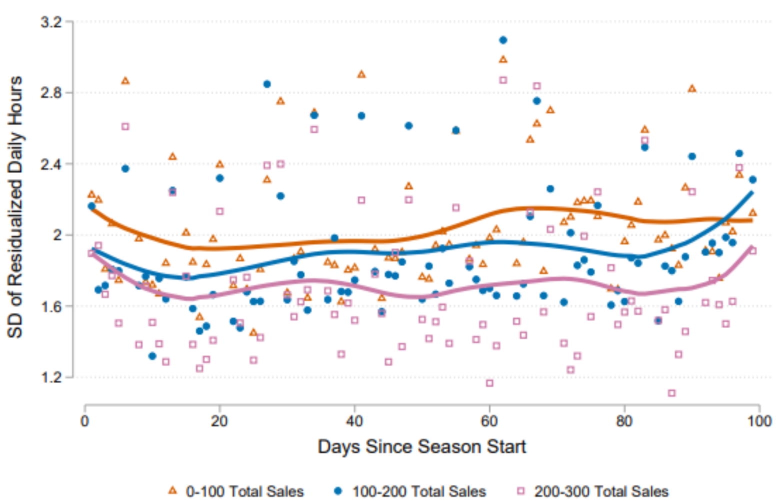

Panel D: Hours Over Progress Toward Total, by Total

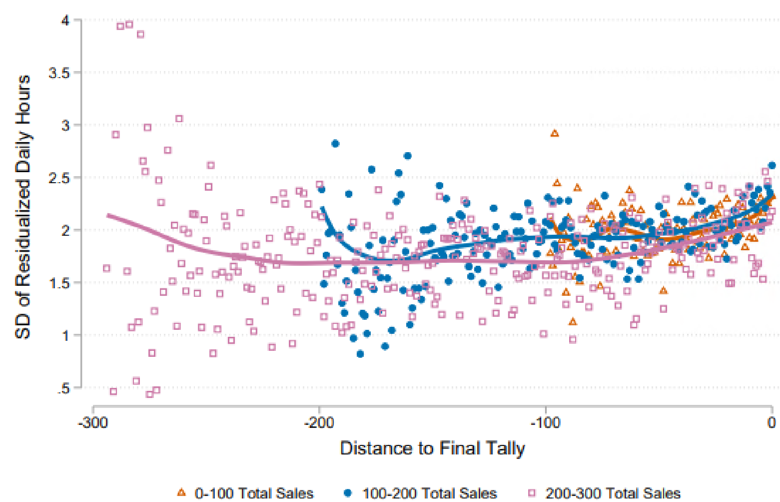

Source: Author's calculations of data from a pest control sales company.

Notes: Residuals come from a regression of daily sales or hours on seller and day-of-the-week fixed effects. Lines represent LOWESS smoothing.

Figure A6: Estimates of Stopping Probability by Month, Experience

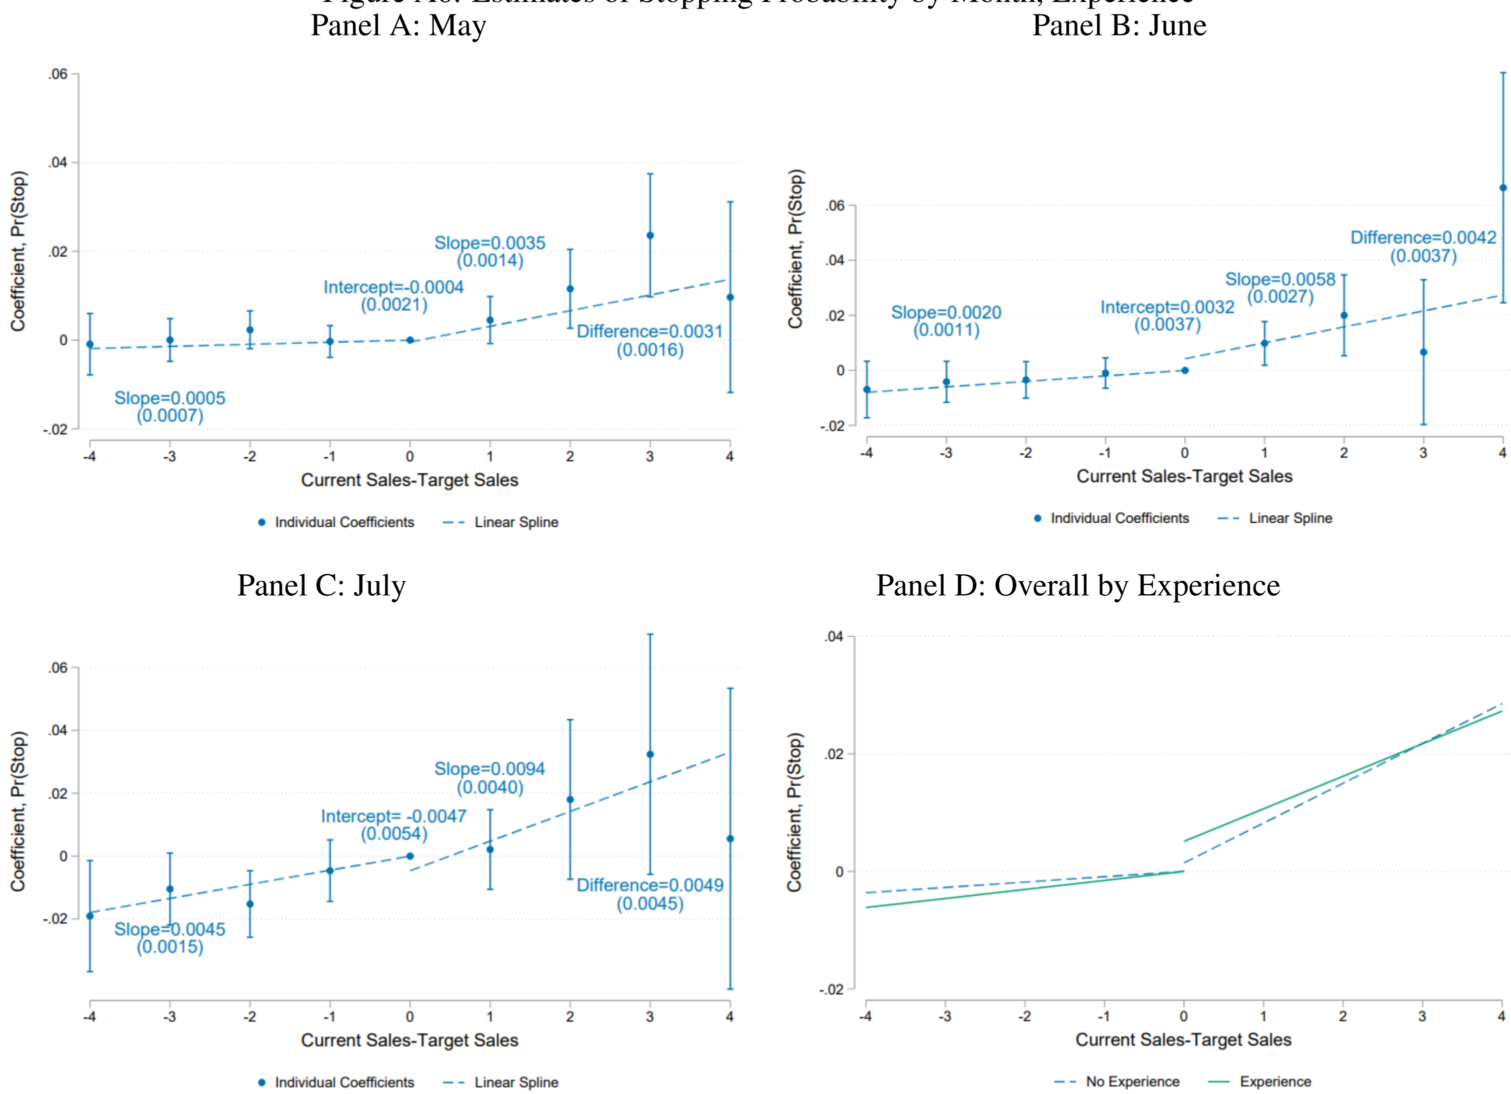

Source: Author's calculations of data from a pest control sales company.

Notes: Results are from estimates of Equations 1 and 2 separated by calendar month. Standard errors are clustered at the seller level. Individual coefficients are rounded to the nearest integer. The coefficient on +1 includes the interval (0,1.5). Panel D shows parametric estimates interacted with an indicator for whether the seller is new to the job ("no experience") or is returning for a second or third sales season in the data ("experienced").

Figure A7: Robustness Test: Pooled Estimates with Tournament/Non-Tournament Interactions
Panel A: Probability of Stopping for the Day

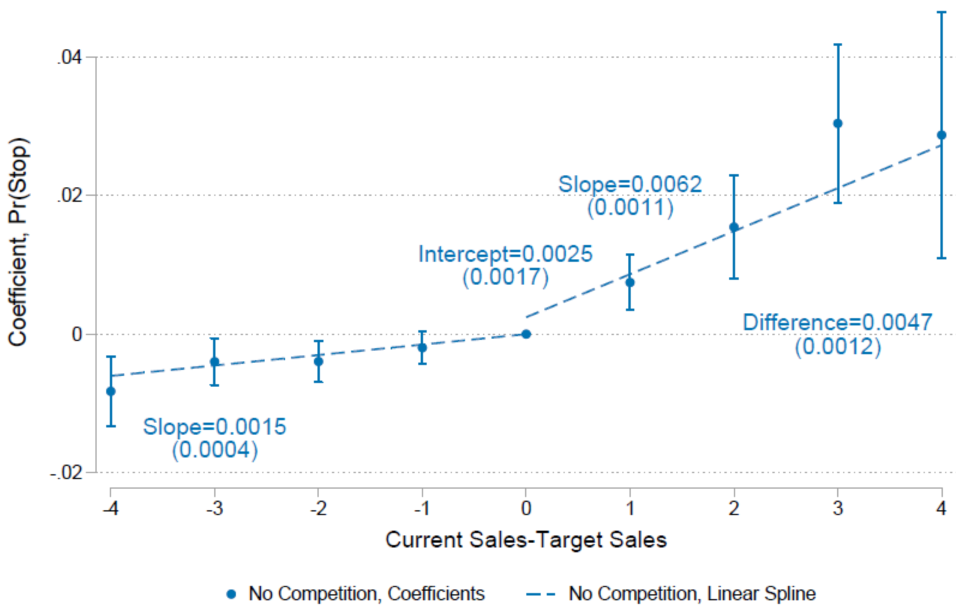

Panel B: Pitches Per Half Hour

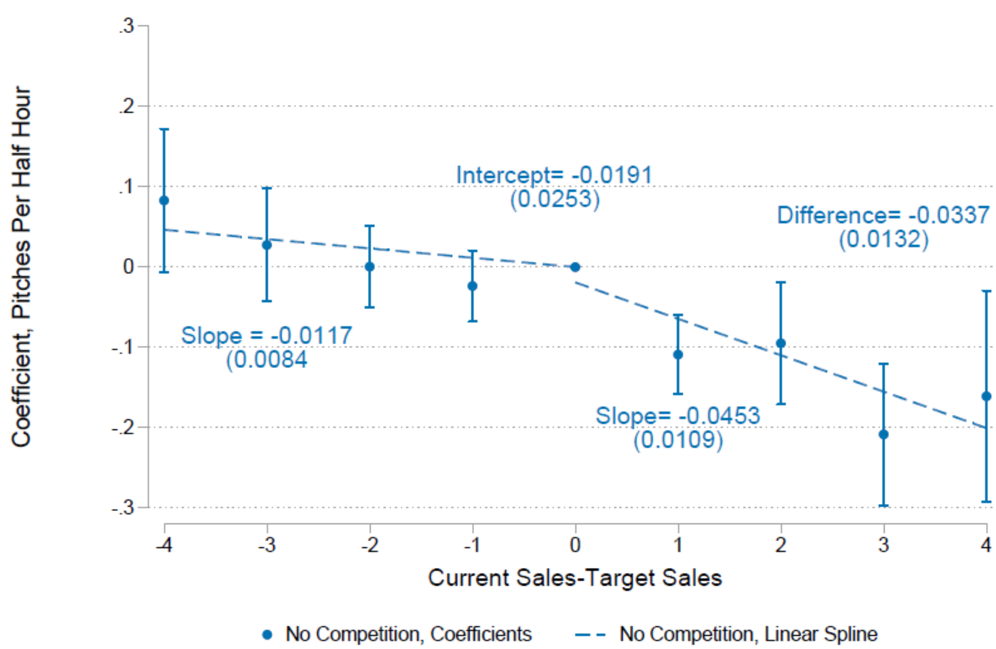

Source: Author's calculations of data from a pest control sales company.
Notes: Results are from estimates of Equations 1 and 2 but pooling all data and including interactions between indicators for tournament type or non-tournament days and distance to the target. Reported results are for non-tournament interactions. Standard errors are clustered at the seller level. Individual coefficients are rounded to the nearest integer. The coefficient on +1 includes the interval (0,1.5).

Figure A8: "Goal-Based" Reference

Panel A: Probability of Stopping for the Day

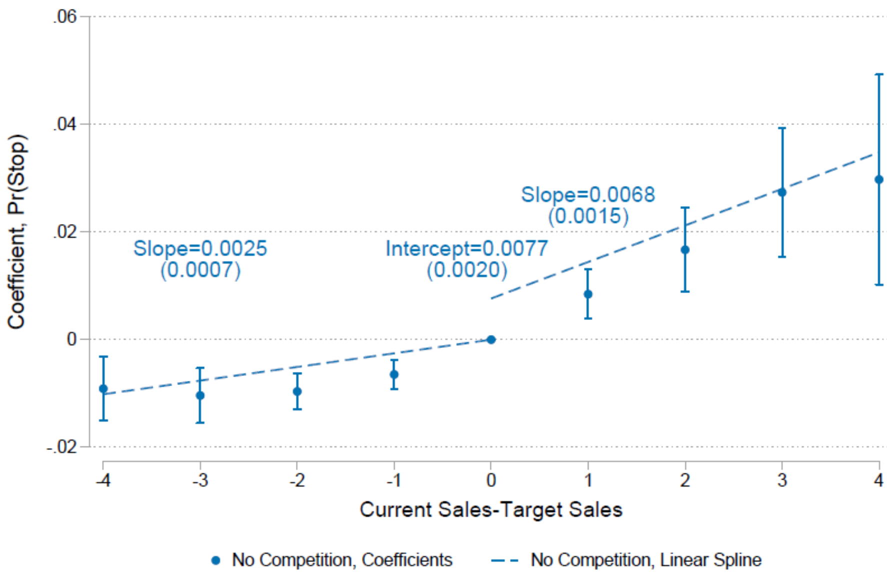

Panel B: Pitches Per Half Hour

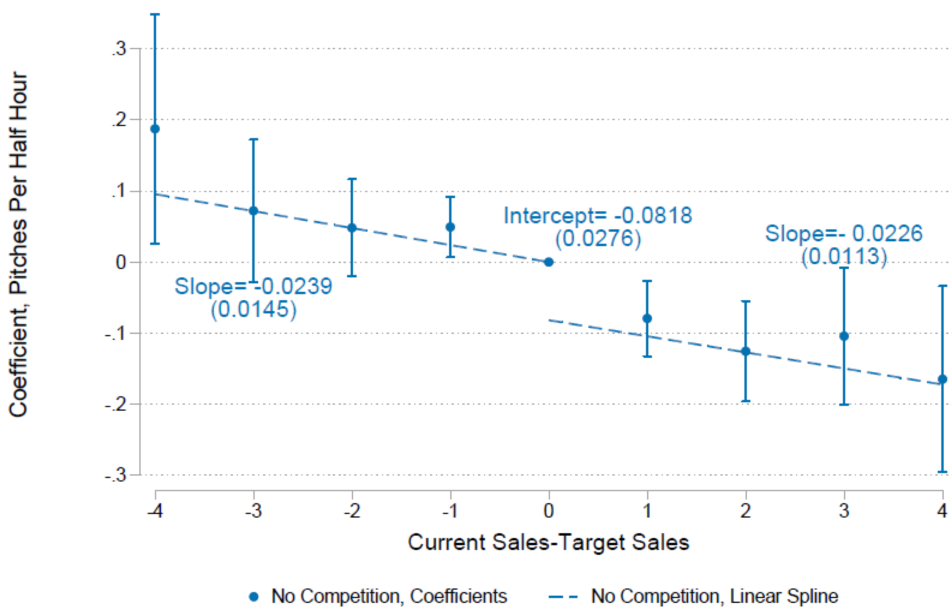

Source: Author's calculations of data from a pest control sales company.

Notes: Results are from estimates of Equations 1 and 2 but pooling all data and including interactions between indicators for tournament type or non-tournament days and distance to the target. Reported results are for non-tournament interactions. The target in these models is a projection of the first two weeks of performance to the nearest bonus threshold at the end of the season. Standard errors are clustered at the seller level. Individual coefficients are rounded to the nearest integer. The coefficient on +1 includes the interval (0,1.5).

Figure A9: Density Estimates of Experimental Performance
Piece Rate vs Bonus at 2,400

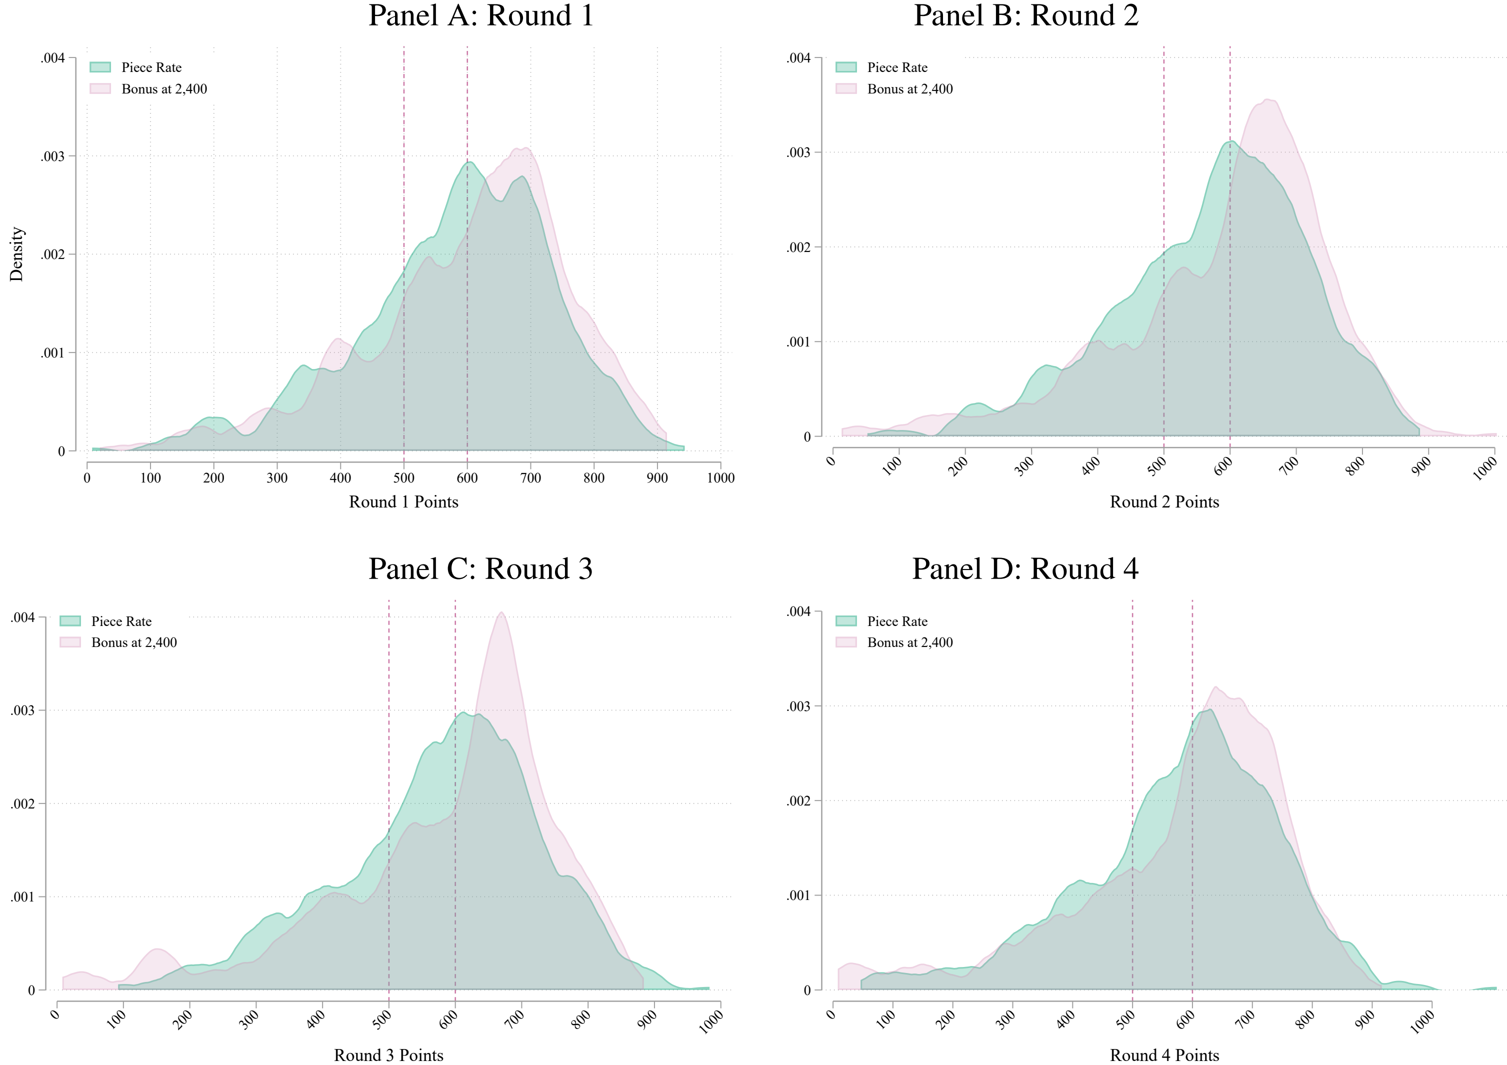

Source: Author's calculations of data from an online experiment.
Notes: These comparisons follow those in Figure 6 comparing the piece rate treatment to the bonus at 2,000 points treatment.

Figure A10: Variance in Each Round, Piece Rate vs Bonus at 2,000
Panel A: Standard Deviation by Round

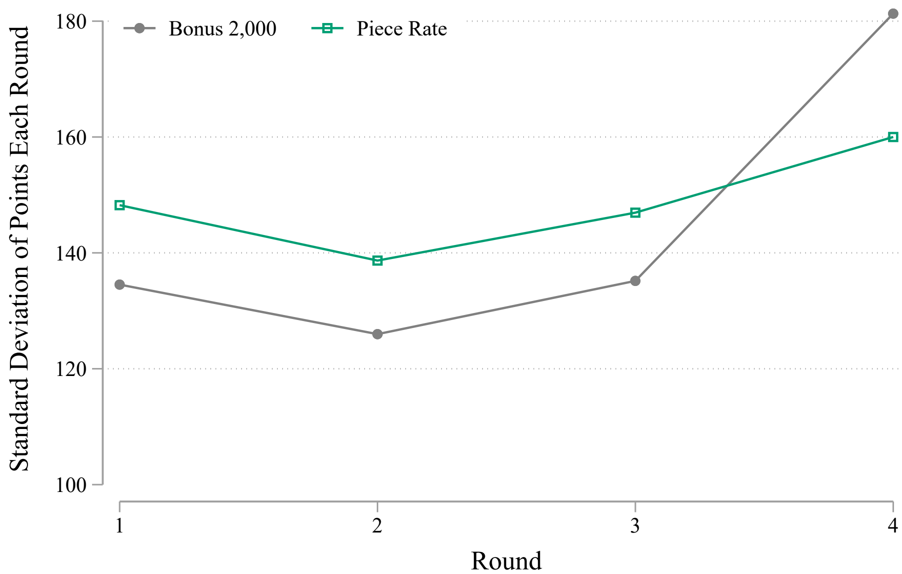

Panel B: Percent Change in Standard Deviation by Round

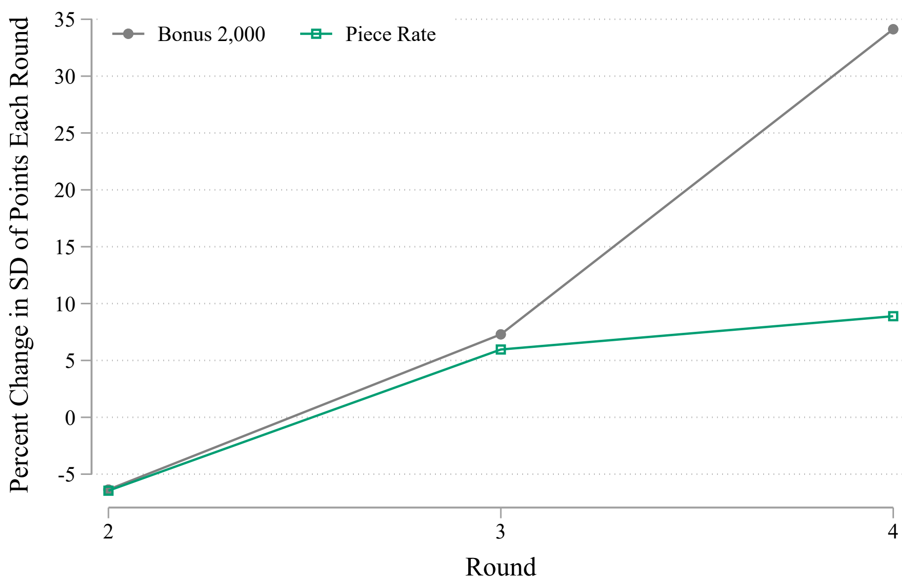

Source: Author's calculations of data from an online experiment.
Notes: Panel B changes are measured relative to the prior round.

Figure A11: Difference in Enjoyment and Stress in Round 1: Bonus vs Piece Rate

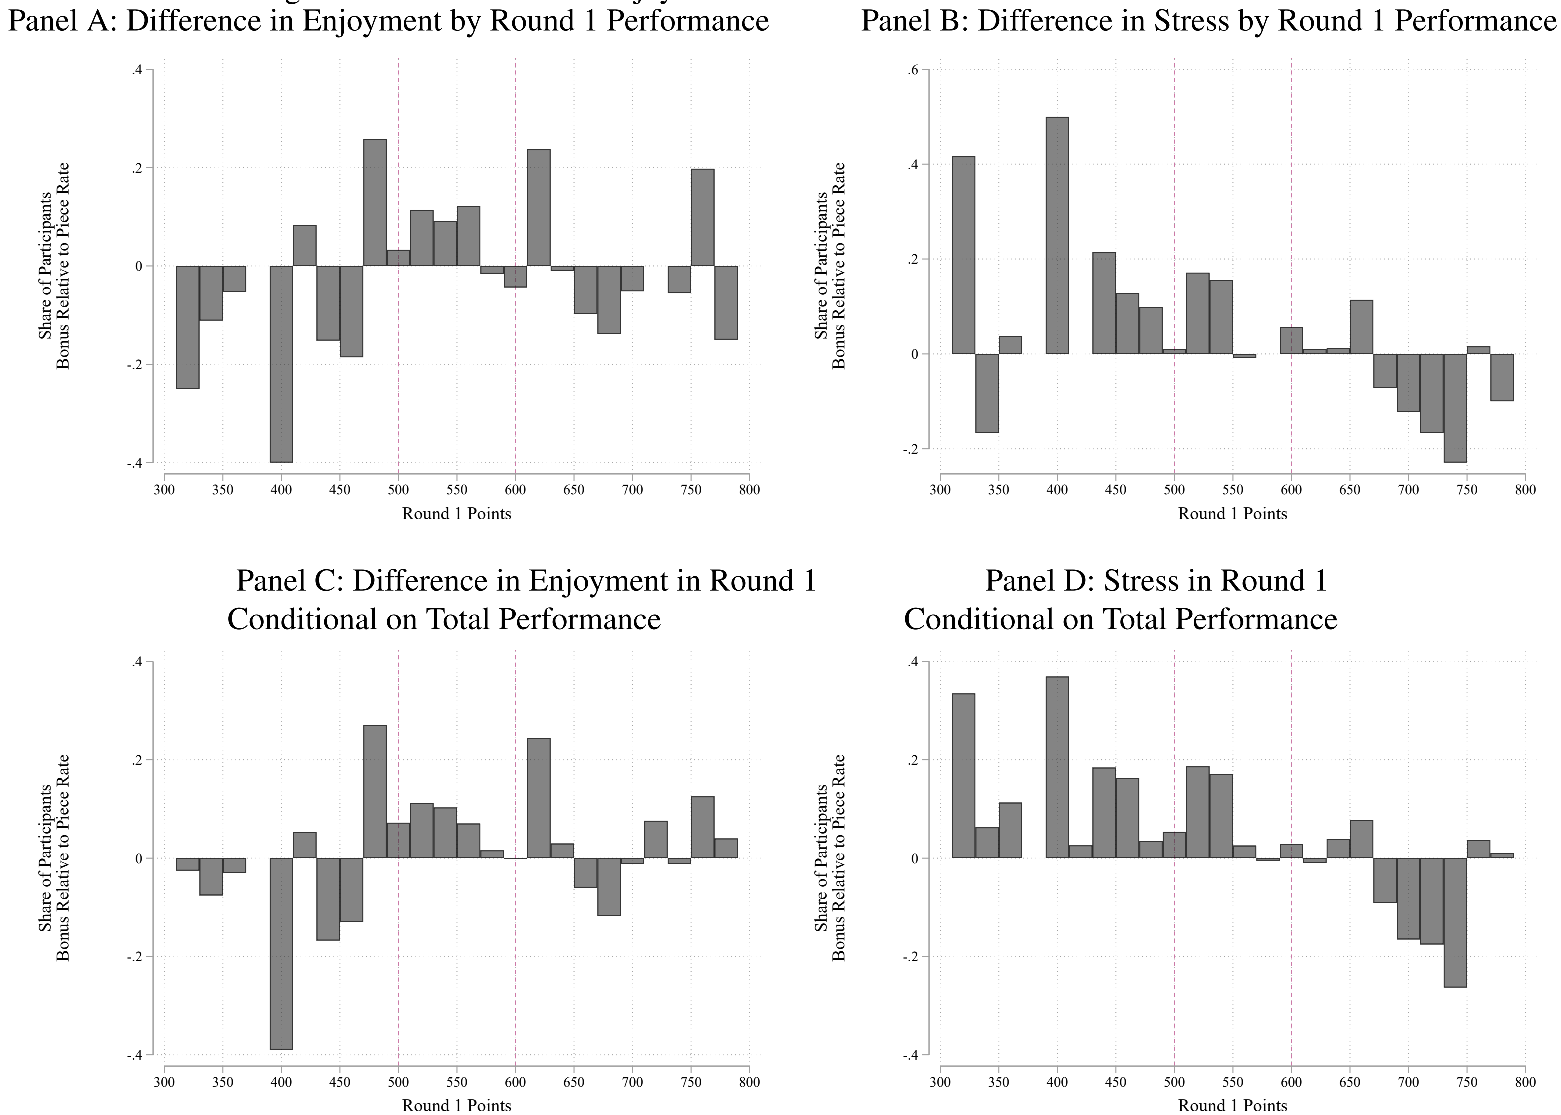

Source: Author's calculations of data from an online experiment.
Notes: The two measures of enjoyment and stress are based on answering "agree" or "strongly agree" that they enjoyed the task or felt stress. Total performance is controlled non-parametrically with bins for every 50 total points at the end of the task (Panels C and D).

Table A1: Summary Statistics of Key Variables

Panel A: Half-Hourly Panel	Mean	SD
Pr(stop)	0.074	0.262
Pitches Per Half Hour	2.281	2.498
Sales Per Half Hour	0.156	0.419

Panel B: Daily Panel	Mean	SD
Sales Per Day	2.02	2.20
Labor Supply		
Pitches Per Day	31.21	19.63
Hours Per Day	6.94	2.23
Average Sales Specific to Day of the Week	1.99	1.60
Weather		
Precipitation (1/10th MM)	4.00	8.52
High Temperature (Celsius)	26.85	5.00
Low Temperature (Celsius)	15.29	4.97
Select ZIP Code Characteristics		
Median HH Income	85,945	25,385
% HH Income $100,000-$150,000	19.49	4.69
% Residents Living in Same Home From Last Year	88.19	4.41
Total Housing Units	112,203	5,766
% Housing Units Single-Family Homes	80.08	11.85
Median Home Value	258,083	107,492
% Non-Hispanic White	80.36	13.71
% Bachelors Degree or More	44.93	14.74
Total Sellers	512	
Total Days	180	
Total Half-Hourly Observations	458,558	
Total Daily Observations	37,984	

Source: Author's calculations of data from a pest control sales company, NOAA daily weather data, and ACS data on ZIP codes.

Table A2: Test of Location Sorting

Sales Per Day, All Significant Coefficients	(1) ACS	(2) Weather	(3) Both
% Non-Hispanic Black	0.0313*		0.0316*
	(0.0161)		(0.0162)
% Single Mothers	-0.0833**		-0.833**
	(0.0403)		(0.0403)
% House Value \$100,000-\$200,000	-0.0276**		-0.0271*
	(0.0140)		(0.0139)
Precipitation (1/10th MM)		-0.00507***	-0.00635***
		(0.00152)	(0.00147)
High Temperature (Celsius)		0.0209**	0.0188**
		(0.00774)	(0.00788)
Low Temperature (Celsius)		-0.0131	-0.0142
		(0.0107)	(0.0108)
Observations	37,508	37,943	37,467
R-squared	0.029	0.013	0.031
F-Statistic	1.59	9.724	3.782
prob>F	0.054	0	0

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Source: Author's calculations of data from a pest control sales company, the American Community Survey 2013-2017 5-year ZIP code estimates, and daily weather data from NOAA.
Notes: Results are from regression of observed ZIP code characteristics from the American Community Survey (ACS) and dailiy weather data on sales generated per day, including day-of-the-week, week-of-the-seaon, and year fixed effects. Standard errors clustered at the seller level. Non-significant coefficients on % Non-Hispanic White, % Hispanic; % of households with income \$50,000-\$75,000, \$100,000-\$150,000, and >\$200,000; median household income, poverty rate, unemployment rate, % adults with Bachelors degree or more, % households in the same home as last year; total housing units, % of housing units that are single-family homes; % homes with value \$100,000-\$200,000, \$200,000-\$300,000, \$300,000-\$500,000, and \$500,000-\$1 million and median home value.

Table A3: Non-Parametric Estimates
Expectations-Based References

Distance to Expectations	(1) Pr(Stop)	(2) Pitches Per Half Hour
-8	-0.0128**	0.174
	(0.00535)	(0.169)
-7	-0.00577	0.00429
	(0.00472)	(0.0765)
-6	-0.00812*	0.0403
	(0.00421)	(0.0888)
-5	-0.0130***	0.120*
	(0.00290)	(0.0688)
-4	-0.00796***	0.112**
	(0.00264)	(0.0469)
-3	-0.00373*	0.0515
	(0.00195)	(0.0360)
-2	-0.00363**	0.0134
	(0.00174)	(0.0277)
-1	-0.00115	-0.0257
	(0.00138)	(0.0242)
1	0.00727***	-0.0939***
	(0.00213)	(0.0258)
2	0.0165***	-0.0659
	(0.00381)	(0.0406)
3	0.0266***	-0.178***
	(0.00586)	(0.0503)
4	0.0263***	-0.114
	(0.00892)	(0.0699)
5	0.00574	-0.224**
	(0.0132)	(0.108)
6	0.00874	-0.244**
	(0.0177)	(0.107)
7	0.0254	-0.420**
	(0.0365)	(0.163)
8	0.0473	-0.209
	(0.0426)	(0.254)
9	0.0704	0.733
	(0.0542)	(0.569)
10	0.199***	-0.819***
	(0.0738)	(0.294)

Robust standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

Notes: Results are from regression in Equation 1 and coincide with estimates from Figure A7. Standard errors clustered at the seller level.

Table A4: Parametric Estimates of Stopping Probability
Pooled Estimates with Interactions for Tournament/Non-Tournament

Panel A: Expectations-Based References

	(1) Slope Below Reference	(2) Slope Change Above Reference	(3) Intercept Shift at Reference	(4) Ratio of Slopes [(Change Above + Below)/Below]
No Competition	0.00151***	0.00470***	0.00244	4.113
	(0.00042)	(0.00124)	(0.00170)	
Individual Competitions	0.000333	0.00003	0.00379	1.090
	(0.00058)	(0.00302)	(0.00399)	
Team Competitions	0.00242***	0.00055	0.00997***	1.227
	(0.00048)	(0.00194)	(0.00227)	
Benchmark Competitions	0.0014**	0.00011	0.00849**	1.079
	(0.00058)	(0.0030)	(0.00376)	

Panel B: Goal-Based References

	(1) Slope Below Reference	(2) Slope Change Above Reference	(3) Intercept Shift at Reference	(4) Ratio of Slopes [(Change Above + Below)/Below]
No Competition	0.00252***	0.00431**	0.00768***	2.710
	(0.00073)	(0.00180)	(0.00206)	
Individual Competitions	0.000238	0.00237	0.00311	10.958
	(0.00099)	(0.00292)	(0.00433)	
Team Competitions	0.00360***	0.00119	0.0029	1.331
	(0.00086)	(0.0022)	(0.00258)	
Benchmark Competitions	0.00316***	-0.00114	0.00207	0.639
	(0.00098)	(0.00271)	(0.00344)	

Robust standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

Notes: Results are from regression in Equation 2 but include interactions between indicators for each tournament/non-tournament period and distance to the reference. Standard errors clustered at the seller level.

Table A5: Robustness Check: Non-Linear Least Squares

Model Parameters	(1) Pr(Stop)	(2) Pitches/Half Hour
Optimal Cutoff	0.11	0
Slope Before Cutoff	0.00074***	-0.0132**
	(0.00026)	(0.0056)
Slope Change After Cutoff	0.0031***	-0.0309***
	(0.0012)	(0.0105)
Intercept Shift at Cutoff	0.0054***	0.0057
	(0.0022)	(0.0270)
Constant	-0.0011***	-0.0097
	(0.0004)	(0.0117)
Ratio of Slopes	5.2	3.3

Robust standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

Source: Author's calculations of data from a pest control sales company.
Notes: Estimates use the residualized outcome variables from a regression on all fixed effects and controls in the non-linear estimates. Standard errors clustered at the seller level.

Table A6: Robustness Check: Parametric Model Adding Exertion Margin as Control

	(1) Pr(Stop)	(2) Pitches/Half Hour
Cumulative Pitches	-0.0005***	0.0261***
	(0.000007)	(0.0013)
Slope Before Cutoff	0.0019***	-0.0175
	(0.0005)	(0.0079)
Slope Change at Cutoff	0.0035***	-0.00005
	(0.0013)	(0.012)
Intercept Shift at Cutoff	0.0007	0.0144
	(0.0018)	(0.0244)
Ratio of Slopes	2.8	0.997

Robust standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

Source: Author's calculations of data from a pest control sales company.
Notes: Results are from estimating Equation 2 but the model includes a control cumulative pitches that day. This adjusts for any effects of fatigue from working more intensely. Standard errors clustered at the seller level.

B Online Appendix: Other Tests of Reference Dependence and Persistence

As an auxiliary battery of tests for reference dependence, I perform two regression exercises to examine performance persistence and the relationship between sales and hours on above- versus below-average work days. The first is a regression of each seller's total sales at the end of the season on their average sales in the first two weeks of the season as well as the first five weeks of the season. A high R-squared indicates that initial daily sales outcomes and labor supply choices have high predictive power for total cumulative sales. Though not conclusive, this indicates a high degree of persistence in performance. In addition, observed persistence in this measure shows that changing one's goal for the final outcome over time does not appear particularly common in the data.

An important question regarding any day-to-day persistence is whether there is scope for workers to adjust their performance or if there are ceiling effects in place. There are two reasons to doubt this explanation for persistence in performance: first, there is significant variation in performance from day to day within person. Conditional on individual seller fixed effects, the standard deviation of performance is 1.7 sales, meaning that a significant amount of variation is still in play, and it is difficult to attribute all of that variation to demand shocks because this would imply that all workers are exhibiting maximum effort already and success is purely defined by local demand conditional on worker skills. Second, during periods in which the sellers are subject to additional incentives (the tournament periods), they significantly increase their performance. In a regression of daily sales on seller, day of the week, week of the season, and year fixed effects with indicators for tournament status, the average seller increases their sales by 0.12 sales (6.3%) *every day* during the tournament period. This provides strong evidence that there is significant room to increase their sales on particular days if presented with different incentives. See Appendix Table **??**.

The second regression is a test of the strength of the relationship between work hours and performance based on exceeding vs not exceeding expectations. I use a panel of worker-day observations to estimate a model of hours worked each day on the number of sales that day interacted with an indicator for if the day's total sales were higher or lower than expectations (average daily sales specific to each day of the week). I estimate:

$$\begin{aligned} y_{idwa} =& \beta_0 + \beta_1 Sales_{idwa} * \boldsymbol{I}^{+}_{Expectations} + \beta_2 Sales_{idwa} * \boldsymbol{I}^{-}_{Expectations} \\ &+ \alpha X_{zdwa} + \sigma W_{zdwa} + \mu_i + \nu_d + \omega_w + \tau_a + \varepsilon_{idwa} \end{aligned} \tag{4}$$

I include fixed effects for seller (μ_i), day of the week (ν_d), week of the season (ω_w), and year (τ_a). The outcome is hours worked that day, while $Sales$ is the total number of service contracts the seller sold that day. The indicators $\boldsymbol{I}^{+}_{Expectations}$ and $\boldsymbol{I}^{-}_{Expectations}$ are dummy variables for if the total sales that day were above expectations or below. In the standard model, because workers will increase their hours when the cost of effort is low, β_1 will be more strongly positive than β_2. In other words, the relationship between work hours and sales will be stronger when total sales for the day are above average (Dellavigna, 2009). The opposite is true under reference dependence.[38]

The results of my two regression exercises using my daily panel are in Table B2. In Panel A, the R-squared for the regression of total sales at the end of the season on average daily sales in weeks 1–2 is 0.752, meaning that average daily performance in the first two weeks explains over three-quarters of the variation in total cumulative sales at the end of the season. Expanding this period to the first five weeks, the R-squared is 0.872, explaining almost 90% of the variation in total sales. There is little unexplained variation in total season sales after conditioning on the first two to five weeks, and there is high congruence between sales outcomes in the first two weeks and behaviors the rest of the season. What this means in practice is that after an initial early learning period, these sellers do not appear to significantly revise their long-term or short-term targets, meaning that once expectations are formed in the first 2-4 weeks of the season, they are extremely stable.

In Panel B, the relationship between daily sales and hours worked is stronger on days that fell *below* expectations compared to days that exceeded expectations. This runs counter to the predictions of the standard model that workers will work more hours on days that have high wage returns.

One might wonder if other external commitments might drive this relationship between sales and hours in Table B2 by putting limits on how many hours a seller can be in the field and suppressing both hours and total sales. The role of these commitments is limited for multiple reasons. The first is contextual: these sales workers are young, mostly unpartnered, and live away from their normal homes, social networks, schools, and other possible employment. The second is statistical. If sellers did have external commitments that would be consistent enough to systematically place these commitments into the "bad day" category (below their mean performance), one might expect these to be correlated with the time of the day, the day of the week, or the week of the sales season. My fixed effects remove variation correlated with these factors. If external commitments are affecting work hours, one might expect these commitments to disproportionately affect workers later in the day (rather than at, for example, 2:00 PM). If external commitments do take sellers out of the field later in the day, that should *strengthen* the relationship between hours and sales on above-average days because, according to the data in Panel C of Figure C1, sales are much easier to secure during later work hours. This would work in the opposite direction from the patterns I find.

[38] That contact and sales rates increase at the end of the work day additionally "stacks the deck" against reference dependence in this model because sales late in the day are less costly to achieve.

Table B1: The Effect of Tournament Incentives on Daily Sales

VARIABLES	(1) Daily Sales
Any Tournament Incentive	0.128***
	(0.0326)
Observations	38,927
R-squared	0.434
Mean Sales	2.034
Percent Effect	6.3 %

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Source: Author's calculations of data from a pest control sales company, NOAA daily weather data, and ACS data on ZIP codes.
Notes: Regression is executed on the panel of daily sales and includes controls for weather and work area ZIP code characteristics. Estimates include fixed effects for seller, day of the week, week of the season, and year. Standard errors clustered at the seller level.

Table B2: Secondary Evidence of Persistence and Reference Dependence

Panel A: Average Daily Sales in Early Weeks

Total Sales at End of Season	Weeks 1–2	Weeks 1–5
Average Daily Sales	95.91***	91.32***
	(3.785)	(2.129)
Observations	33,728	36,857
R-squared	0.752	0.872

Panel B: Sales and Hours, Days that Exceeded Expectations or Not

Hours Worked Per Day	Did Not Exceed	Exceeded Expectations
Sales	0.441***	0.335***
	(0.0179)	(0.0101)
Observations	37,977	
R-squared	0.266	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Source: Author's calculations of data from a pest control sales company.
Notes: Panel A is from a regression of sellers' total sales at the end of the season on average daily sales during the first two or five weeks of the season. Panel B is from Equation 4 and includes fixed effects for seller, day of the week, week of the season, and year. Standard errors are clustered at the seller level.

C Online Appendix: Data

The pest control sales company data were obtained through a data use agreement prohibiting disclosure of the company's identity or intimate details of their operations.

The data cover the entirety of all sales and knocks recorded from January 2018 to January 2020. Sales in the "off-season" are not compensated the same way as they are during the summer, and there are very few recorded knocks in their system. Most sales the company generates during the off-season are renewals of current contracts for the following year as well as follow-ups with past customers, but those contacts are typically not done in person. Most off-season knocks are those done in the service of training new sellers. The knocking data are reported using their common application, which also shows leaderboards, team performance, and the performance of all other sellers in the company. The centralized sales website also contains sales information but does not include knocking information. Competition rules, dates, and prizes were collected from raw internal company documents as well as the company website usually available only to contractors and employees.

To correctly measure the incentives and behavior of these workers at the right time, I impose a few basic restrictions to my half-hourly panel. I limit my sample to the "summer sales season" each year, which is the period from the last week of April to the third week of August. This excludes trainees who arrive early, those who stay late into the end of August or early September (who are usually managers and those not enrolled in school), and off-season sales. I exclude the last two weeks of August because participation drops precipitously as sellers return to school. Less than 50% of sellers stay past August 17th-18th, and less than 25% of sellers stay past August 25th-26th. I then exclude any sellers who stopped working altogether before late May, which effectively excludes the least able sellers who averaged less than one sale per week and decided to go home after experiencing this lack of success. This group also includes managers who record knocks for training purposes during the first month. Off-season sales during the September to April months entail a different compensation structure, and many of the sales are generated by full-time employees of the company rather than the independent contractors that work during the summer.

In my half-hourly panel, I exclude observations with no previous expectations, i.e. the first week a seller is active. In all, my half-hourly panel consists of approximately 459,000 observations for 512 sellers across 180 days in 2018-2019.

From an incentive standpoint, if there is a positive autocorrelation in sales each day—that is, if success now is predictive of success in the near future—then a worker having success right now faces lower marginal costs of effort in the coming hours. This will work against the downward shift in labor supply predicted by Prospect Theory. To test for this, I residualize sales each half hour by regressing sales each half hour on fixed effects for seller, day of the week, week of the season, and year as well as controls for actively knocking on doors, weather, and ZIP code characteristics. I then calculate the autocorrelation in these residuals between half-hour periods and present the results in Panel A of Figure C1. The results suggest that there is some autocorrelation in residualized sales for just under one hour, or that success now is predictive of success at least for the next half hour.

This relationship appears driven by the positive relationship between the time of day and sales after 5:30 PM. After conditioning on the length of the shift and time of day, Panel B of Figure C1 shows little autocorrelation. Panel C of Figure C1 shows that average seller performance increases as the day progresses, particularly after 5:30 PM when residents return home from work. This is not due to a change in the composition of workers, but because workers have more opportunities to make contact with residents. The marginal cost of sales falls later in the day. If sellers understand these dynamics, they have an incentive to continue to work. In short, despite these contextual features "stacking the deck" against reference dependence in terms of incentives, I still detect meaningful evidence using my formal tests, which I describe in Section 5.1.

Figure C1: Upward Pressures on Labor Supply During the Day

Panel A: Autocorrelation of Sales

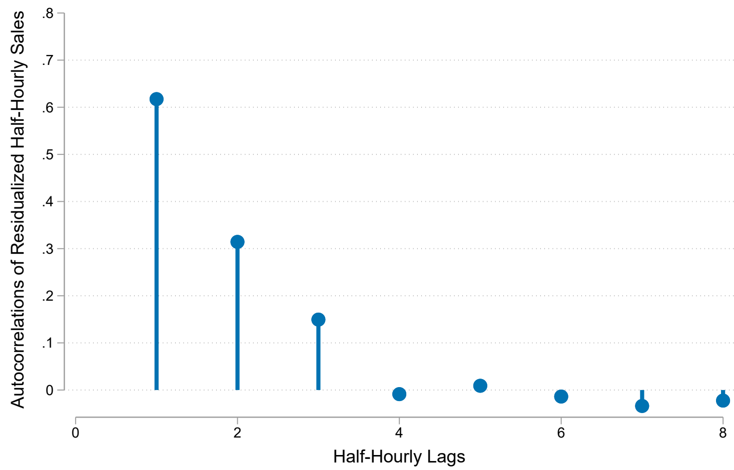

Panel B: Autocorrelation of Sales Conditional on Time of Day and Shift Length

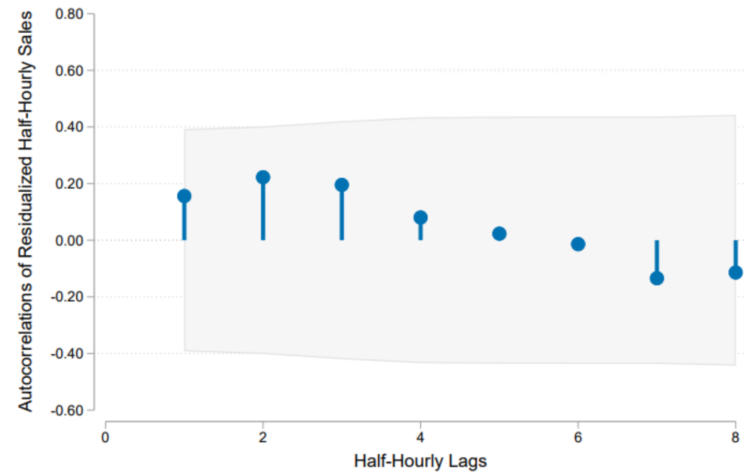

Panel C: Mean Sales by Half Hour

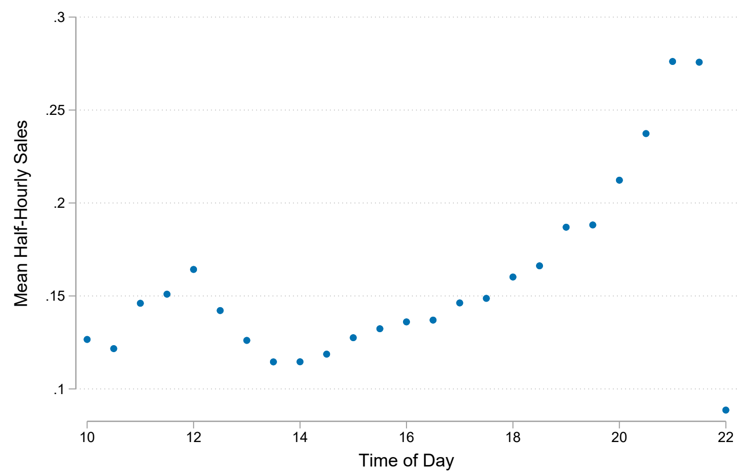

Source: Author's calculations of data from a pest control sales company.

Notes: In Panel A, residualized sales come from a regression of sales each half hour on seller, half-hour-of-the-day, day-of-the-week, week-of-season, and year fixed effects as well as controls for having any knocks recorded that half hour, weather, and ZIP code characteristics. I then calculate the autocorrelation for those predicted residuals for half hour lags of one through eight. Only the correlation for the one-period lag is statistically significant. Panel B adds additional controls for half hour of the shift and half hour of the day. The shaded region shows Bartlett's formula for MA(q) 95% confidence bands.

Figure C2: Autocorrelation in Daily Sales

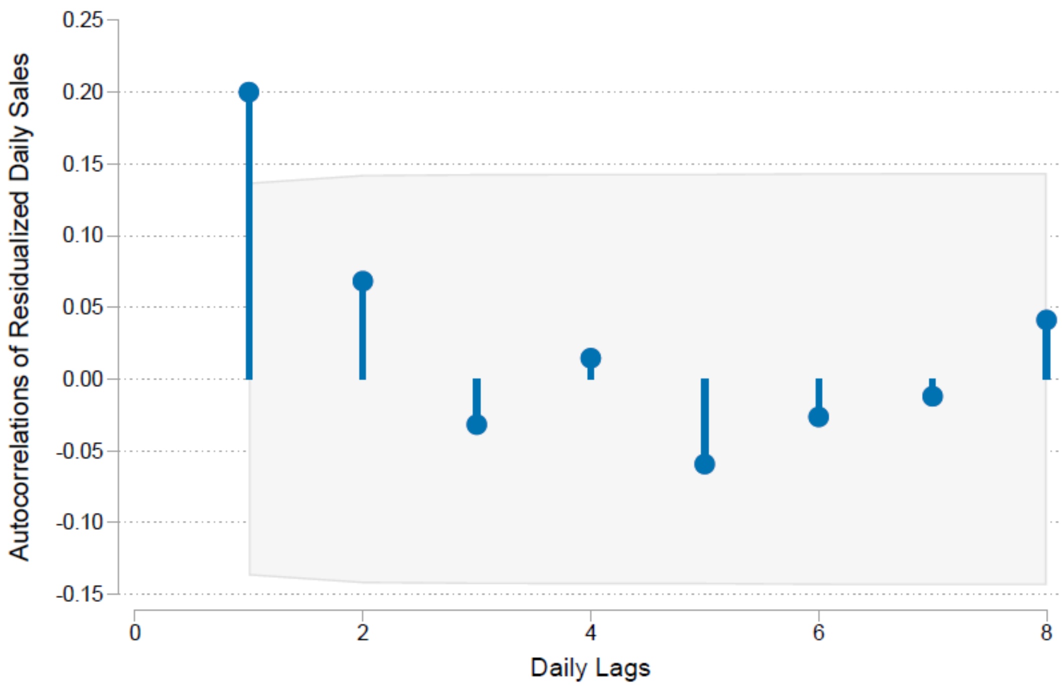

Source: Author's calculations of data from a pest control sales company.
Notes: This figure uses the seller-day panel to calculate residualized sales. I regression of sales each day on seller, day-of-the-week, week-of-the-season, and year fixed effects as well as controls for having any knocks recorded that day, weather, and ZIP code characteristics. I then calculate the autocorrelation for those predicted residuals for lags of one through eight days. The shaded region shows Bartlett's formula for MA(q) 95% confidence bands. The low autocorrelation between days indicates that performance today is not strongly predictive of performance tomorrow, or that individual workdays come from independent draws.

D Online Appendix: Further Background

The company whose data I use (which I call "PestCo") operates a full-service pest control service operation. In addition to removing insects, spiders, and rodents, they apply preventative treatments to prevent pests from returning or growing larvae near an individual home. There is a range of services they provide, and sellers are encouraged to "upsell" for more comprehensive services whenever they see an opportunity. Sellers are given the responsibility to generate new contracts and schedule the service with a separate wing of the company that performs the service. Most contracts last 12–18 months. Commission rates are based on the annualized value of the contracts the seller generates.

PestCo is not markedly different from the rest of the sales industry in terms of its use of incentive schemes. Their independent contractor agreements and practices are all in line with industry standards.

Sellers are paid an up-front portion ($75) of their commissions during the two-week period each sale is made, similar to a regular paycheck. The balance of commission payments is calculated at the end of the season after the status of all contracts is known. Final payouts for Spring sales are given in the Fall, and late Summer sales payouts are given at the end of the year. Most contractor agreements include penalties for leaving the selling area before the official end of the sales season or for not recording knocking activity a minimum number of days. The penalties typically stipulate that regardless of the number of sales, the commission the seller earns will return to some low base rate (usually 18–20%).

Prior to leaving for their assigned metro area, sellers at PestCo are trained in sales techniques and are given a detailed manual of behavioral tools to help them over the course of the summer. These include training on proper body language, handshaking, standards for appearance, overcoming customer objections, rephrasing customer concerns, interacting with upset neighbors, and how to look for and identify pests before approaching a door. They are provided with video examples of strong sales performance and are encouraged to review their training materials on a daily basis.

PestCo takes an active role in trying to motivate their workers. In training materials, the company encourages their sellers to be physically active and healthy, to be honest about their performance and goals, and to take accountability for their own performance and summer experience. These training materials are especially important because approximately half of the sellers who are working any given day are brand new to the company and the industry. Sellers are encouraged to learn advanced sales techniques from their more experienced teammates or roommates.

Work neighborhoods for each seller are assigned by a local team leader. Metro areas are divided into sections for each team, and within their section, team leaders assign sellers to a neighborhood. Work in each neighborhood continues until approximately 75% of doors have been marked in their tracking software, after which the seller can request a new area. Managers insist that "work area does not matter" in their training materials, and the evidence I present supports this argument. Area assignments, while not strictly random, are not correlated with sales outcomes in any meaningful way either across or within seller (see A2) and Section 4. Managers emphasize that making assignments to work areas based on perceived skill or other seller attributes is costly to them as managers and generates unclear returns, which undermines the business case for sorting. For example, assigning a better seller to a "harder" neighborhood may generate sales that would otherwise not take place. However, the marginal cost of achieving those may be high and the benefits may be smaller than the difference in sales *speed* across sellers in an "easier" neigbhorhood. In practice, managers do not typically spend large amounts of time on these assignments.

In addition to the high-powered cash incentives built into their contractor agreements, PestCo also runs frequent short-run tournaments for prizes valued from \$300 to \$3,000. These take three forms: individual rank-order, team rank-order, and what I call "benchmark" competitions. Individual rank-order tournaments pit sellers against each other for a single day, and the seller with the most head-to-head daily "wins" at the end of the two-week tournament gets a prize.

Team rank-order tournaments have a similar structure but are based on wins against another team, and "wins" are based on per-seller team revenue. During "benchmark" competitions, if a seller generates more revenue during the week-long competition period than he did during any prior week in the season, he will get a prize. Prizes include merchandise like Bluetooth headphones, apparel, and expensive grills as well as "experiences" like a cruise, resort stay, or annual ski passes, though sellers have the option to cash out the value of the prize. Importantly, these benchmark tournaments occur later in the season, and the prizes are not valuable enough that it is worthwhile for sellers to lower their effort earlier in the season in the hopes of making it easier to gain these prizes. These tournaments are not the focus this study. However, they are important for contextualizing my empirical models because they modify the incentive structure within particular sales days and thus may shift a worker's expectations on that particular day. See Table B1. However, there is no evidence of a more permanent change in expectations.

E Reference Dependence with Loss Aversion and Recent Expectations

The basic insight of models of reference dependence and loss aversion propose that losses loom larger than gains.

O'Donoghue and Sprenger (2018) present a simple model of this idea that is instructive. A worker can choose an effort level e, which yields output $x(e)$ and has a cost of effort $c(e)$. The function $c(e)$ is increasing and convex. Utility is linear in $x(e)$. Suppose there is an output or income reference, r, which can be endogenously determined by rational expectations or exogenously imposed. Distance from the reference, $x(e) - r$, enters the utility or value function:

$$U(e) \equiv x(e) + \mu\left(x(e) - r\right) - c(e) \tag{5}$$

where

$$\mu(z) = \begin{cases} \eta z & \text{if } z \geq 0 \\ \eta\lambda z & \text{if } z \leq 0 \end{cases}$$

The μ function captures "gain-loss utility." The equilibrium labor supply for this utility function with gain/loss utility is given by:

$$\begin{array}{ll} (1+\eta)x'(e) - c'(e) = 0 & \text{if } x(e) - r > 0 \\ (1+\lambda\eta)x'(e) - c'(e) = 0 & \text{if } x(e) - r < 0 \end{array} \tag{6}$$

The shift across the reference threshold reflects the difference in the marginal value of income. At the same level e, the marginal benefit on the left side of the reference ($x(e) < r$) is scaled by a factor of $\lambda > 1$ relative to the right side of the reference ($x(e) > r$). This parameter is the coefficient of loss aversion. The parameter η is the weight of gain-loss utility in the utility function. This simple model with linear utility implies that, if current earnings, $x(e)$, are below the reference, equilibrium labor supply will be higher than if earnings are above the reference for the same value of e. For a loss-averse worker, upon reaching the reference, r, there is a downward kink in the marginal value of income, so labor supply will also kink downward, holding constant effort costs at $c(e)$. Figure E1 shows an illustration of this concept. The marginal utility when $\lambda = 1$ is the same on either side of the reference. However, when $\lambda > 1$ and income is below the reference, the marginal utility is higher and overall utility is lower because being below the reference creates a sense of loss. In the standard case in Equation 6, $\lambda = 1$ or $\eta = 0$, and there is no discontinuous change in marginal benefit across the reference.

The prior literature on labor supply has almost exclusively focused on daily references. This focus simplifies the theoretical tests of reference dependence by limiting the role of income effects, which standard theory suggests may be notable in the long-run but will be negligible each day because daily income plays such a small role in long-run or lifetime earnings (O'Donoghue and Sprenger, 2018; Dellavigna, 2009). This justifies the use of linear utility in Equation 5. Reference dependence with loss aversion predicts in my context that when a seller surpasses her daily reference, the probability she stops working for the day will kink upward, holding other factors constant.

On the other hand, the standard model predicts that if the wage return, $x'(e)$, shifted upward for the same value of e, the worker would unambiguously work more hours regardless of which side of r she is on. When daily wages are high, the standard worker will increase daily labor supply, and when daily wages are low, the worker will stop working earlier in the day. These labor supply decisions will be a smooth function of $x(e)$ and $c(e)$.

As important as the parameter of loss aversion (λ) is to the model of reference dependence, equally important is the definition (or location) of the reference point itself. In an essential theoretical paper (Kőszegi and Rabin, 2006), the KR model theorizes that "recent expectations" act as important reference points. But how do people form short-term expectations? The KR model proposes that these expectations are determined in what they call "personal equilibrium," that is, by behaviors that are optimal given the worker's expectations about the future. Put another way, a forward-looking worker can make a plan around what she perceives to be the optimal path forward, and when the final choice is made in real-time, the planned path becomes her reference point. This "path" is her personal equilibrium. Kőszegi and Rabin (2006) posit when introducing this theory that firms can play a significant role in establishing a worker's personal equilibrium; however, empirical evidence for this role is generally sparse.

This theoretical result has important implications. The first is that if wage increases are anticipated or predictable, a worker will respond by planning to work more hours, or by adjusting her adjusts her planned path, similar to the standard model. In the context of the bonuses paid in the door-to-door sales setting and in my experiment, this means that workers make their initial daily labor supply choices based on what they determine to be optimal given what they expect to be their ability to reach a bonus threshold. If workers obtain new information about their abilities, they can quickly adjust their future goals and their short-term reference points. This creates a feedback loop between future expectations and recent experience wherein a simple measure of average past performance integrates both pieces of information. The second key implication is that workers exhibit gain-loss utility over outcomes that deviate from expectations (the path). After setting her plan for the path ahead, the worker responds each period to whether her performance is below or above what she expects for that period. Significantly, negative comparison utility and higher marginal utility while working below short-term expectations induce more effort.

E.1 Reference Dependence and Goal Setting

In addition to expectations service as reference points, a person's own goals may serve as reference points and can be translated from the long run to the short run.

It is worth explicitly exploring this dynamic of goal-setting under the framework in Koch and Nafziger (2020). Suppose workers perform the same task each period or day (in time $t \in [1, T]$) with effort level e_t that incurs costs $c(e)$ that are convex. Then suppose there is a total benefit b at the end of a long-run evaluation period that is a function of total effort, and effort is deterministic over utility outcomes. If a worker is a quasi-hyperbolic discounter (Laibson, 1997), then there are t versions of the worker, one for each day, with utility $U_t = u_t + \beta[\sum_{\tau=t+1}^{T+1} u_\tau]$ and instantaneous utility u_t and a present-bias factor of β. Instantaneous utility is $u_t = -c(e_t)$, and final period utility $u_{T+1} = \sum_{t=1}^{T} b(e_t)$. Ex ante, a period 0 self sets marginal costs and benefits equal such that $\beta = 1$ and $b'(e_0^*) = c'(e_0^*)$. This would be the equilibrium effort under the standard model of some chosen long-run outcome.

Now suppose each period's self after period 0 discounts future benefits by $\beta < 1$. Equilibrium effort with present-biased preferences would be $\beta b'(e_0^*) = c'(e_0^*)$. A worker who set out to perform at e_0^* to achieve total benefit $b(e_0^*)$ in time 0 has an incentive in time t to substitute effort from today to tomorrow or from the current period to the next. The prospect of substituting effort across days (because total outcomes are fungible across days) may lead to suboptimal effort in time t under the ex-ante assumption that the worker may increase effort in $t + 1$. Importantly, if the benefit at the end of the period (for example, a total payout for a worker) were increased by some proportion γ, while the worker has an increased incentive to gain benefit $\gamma b(e_0^*)$ at the end, the utility benefit each period would only increase by $\beta * \gamma b'$. Practically speaking, that means present bias blunts the incentive effect of additional benefits to perform in the longer run, making such incentives less cost-effective.[39]

But suppose self 0—a forward-looking agent—sets a narrow bracket through a daily or period-specific goal to bind the incentives for self t in the future through additional comparison utility penalties, i.e. for $e_t < g_t$, $\hat{\beta}(g_t - e_t)$. For a sophisticated individual who correctly predicts β and calibrates $\hat{\beta}$, personal equilibrium suggests that g_t should be the same as the optimal effort that period 0 self would choose given their beliefs about future effort, or in other words, that $\hat{e}_{t,0} = g_t$. When tasks are repeated daily, $g_t = [b(e_0^*)]/T$. Self t then provides effort g_t each period, thus solving the self-control problem.

Combating suboptimal effort substitution is the key incentive introduced by narrow bracketing. In the case of daily or period-specific goal-setting, because the marginal utility of income is higher in the loss domain, workers have the incentive to exert more effort on "high-cost" days to achieve a minimum performance. On "low-cost" days, they surpass their target more easily, but the marginal

[39]In in many occupations like sales, effort costs to achieve the same objective may fluctuate through a day-specific cost function ($c_t(e_t)$), that is, the time and effort cost of achieving the same objective. The standard model predicts a worker will provide more effort on exogenously "good days" where the marginal costs of effort are low and less on exogenously "bad days" where the marginal effort costs are high, i.e. when $c_t'(e_t)$ is high. That is, when the marginal benefits are consistent from day to day, higher marginal costs will lead to lower equilibrium effort. Effort, therefore, will fluctuate from day to day. When there is present bias, beyond just a β discount, a worker has further incentive on "bad days" to implement effort substitution because of the expectation of future "good days" to make up for it. This is an implication of Proposition 2 in Koch and Nafziger (2016).

utility of additional income falls, so the worker has the incentive to reduce their labor supply upon surpassing it. Thus, for the worker, there is a cost to narrow bracketing: negative comparisons in the loss domain reduce experienced utility while in that domain. Therefore, workers using narrow brackets as commitment devices will do so only until the broader goal is reached, after which there is no reason to continue engaging in negative comparisons. Because of these costs and despite the possible positive effects on goal attainment, not all workers may engage in this behavior.

From a firm's perspective (or any other principal in a principal-agent setting) where the worker reaching a certain level of output matters most, it is advantageous to induce workers to engage in this dynamic targeting behavior if narrow bracketing leads to higher rates of goal attainment, which might be expected if present bias is common. Thus, two important open questions are: 1) is there evidence that this type of short-run reference setting occurs in labor markets as a commitment device? And 2) can firms induce such behavior; and if so, how, and what are the consequences? In the sales data and experiment, I do not directly test for present bias but rather if workers behave and plan *as if they expect* present bias.

Figure E1: Illustration of Basic Model of Reference Dependence with Loss Aversion

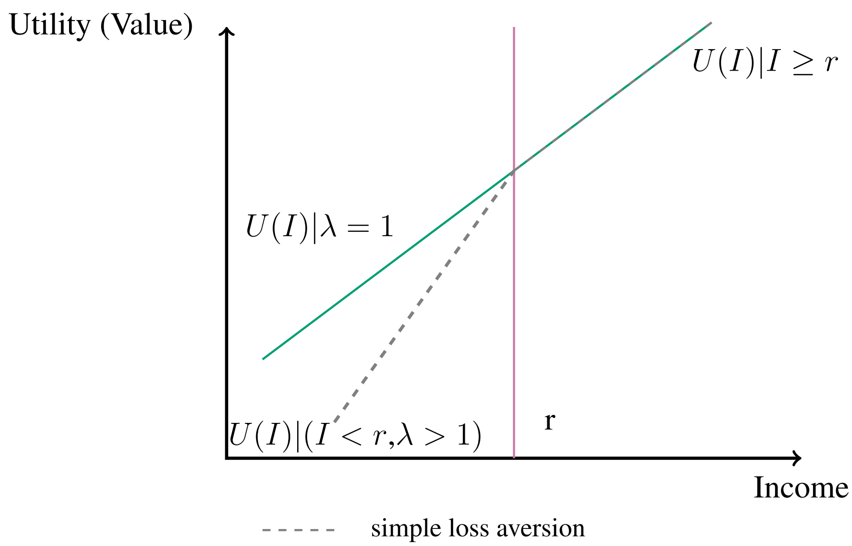

Notes: Illustration of basic loss aversion with linear utility over income. When $\lambda = 1$, the marginal utility above the reference r is the same as marginal utility below the reference.

F Online Experiment Protocol (For Online Publication)

Experimental Protocol

Start of Block: Consent

Q1
Introduction

Welcome to this research project! We very much appreciate your participation.

This is a study about decision-making. Several research institutions have provided funds for this research. It is very important for the success of our research that you answer honestly and read the questions very carefully before answering.

Procedures
You are given instructions on your screen before every decision. Please always make sure to read the instructions carefully before you continue.

Payment
Your payment will consist of the participation fee plus the amount of bonus points that you accumulate throughout the study. The exact amount of bonus points that you receive will depend on your and/or others' decision.

Your bonus will be paid to you using the bonus system within a few days after completion. Your payment for taking part in the study will be sent to you shortly after submission.

Participation
Participation in this research study is completely voluntary. You have the right to withdraw at anytime or refuse to participate entirely without jeopardy to future participation in other studies conducted by us.

Confidentiality

All data obtained from you will be kept confidential and stored on a GDPR compliant secure server. Your prolific ID will only be used for payment purposes. Any personal information that could identify you will be removed or changed before files are shared with other researchers or results are made public in open science repositories.

Verification

At the end of this survey, you will be automatically redirected to Prolific.

Questions about the Research
If you have questions regarding this study, you may contact
thechoicelab@nhh.no

Q2 If you have read and understood the instructions above and want to participate in this study, write ACCEPT in the box below.

____

End of Block: Consent

Start of Block: Attention Check

Q3 Please indicate your agreement with the following statement below:

	Strongly Disagree (1)	Disagree (0)	Agree (0)	Strongly Agree (0)
I need to cover myself in lava so that I don't freeze at night. (1)	○	○	○	○

End of Block: Attention Check

Start of Block: Treatment 1 – Bonus at 2,000 Points

Q4 As part of this study you will play a simple button-pressing task. The object of this task is to alternately press the 'a' and 'b' buttons on your keyboard as quickly as possible. You will play for a **total of 10 minutes**.

Every time you successfully press the 'a' and then the 'b' button, you will receive a point. Note that points will only be rewarded when you **alternate button pushes**: just pressing the 'a' or 'b' button without alternating between the two will not result in points. Buttons must be pressed by

hand only (key-bindings or automated button-pushing programs/scripts cannot be used) or the task will not be approved.
Feel free to score as many points as you can.

<u>You will be paid an extra 1 dollar if you score at least 2,000 points.</u> We will divide your 10-minute playing time into **4 rounds.** This means you will need to score an average of at **least 500 points per round** to receive the bonus.

You will have a 10-second break in between each round. On the next page is an example of how the task will work. Try pressing 'a' and 'b' alternately to score points. We have limited the point total below to a maximum of 50 as this is just practice, but the actual task will not have a limit. Try to reach 50 **<u>as quickly as possible.</u>**

Page Break

JS

Q6 Press 'a' then 'b'...

Points: 0

Page Break

Q7 **<u>You will be paid for your performance.</u>** Please confirm you understand this.

○ Yes (1)

○ No (2)

Page Break

Q8 Proceed to the next page when you are ready to play the task. Your 10-minute task will begin immediately when the page loads and will be divided into 4 rounds.

Page Break

Button-
Pushing Task

End of Block: Treatment 3 – Bonus at 2,400 Points

Start of Block: Treatment 2 – Bonus at 2,400 Points

Q23 As part of this study you will play a simple button-pressing task. The object of this task is to alternately press the 'a' and 'b' buttons on your keyboard as quickly as possible. You will play for a **total of 10 minutes**.

Every time you successfully press the 'a' and then the 'b' button, you will receive a point. Note that points will only be rewarded when you **alternate button pushes**: just pressing the 'a' or 'b' button without alternating between the two will not result in points. Buttons must be pressed by hand only (key-bindings or automated button-pushing programs/scripts cannot be used) or the task will not be approved.
Feel free to score as many points as you can.

As a bonus, you will be paid an extra 1 dollar if you score at least 2,400 points. We will divide your 10-minute playing time into **4 rounds.** This means you will need to score an average of at **least 600 points per round** to receive the bonus.
You will have a 10-second break in between each round. On the next page is an example of how the task will work. Try pressing 'a' and 'b' alternately to score points. We have limited the point total below to a maximum of 50 as this is just practice, but the actual task will not have a limit. Try to reach 50 **as quickly as possible.**

Page Break

JS

Q25 Press 'a' then 'b'...

Points: 0

Page Break

Q26 **<u>You will be paid for your performance.</u>** Please confirm you understand this.

- ○ Yes (1)
- ○ No (2)

Page Break

Q27 Proceed to the next page when you are ready to play the task. Your 10-minute task will begin immediately when the page loads and will be divided into 4 rounds.

Page Break

<u>Button-pushing task</u>

End of Block: Treatment 3 – Bonus at 2,400 Points

Start of Block: Treatment 9 – Piece Rate

Q42 Shortly, you will play a simple button-pressing task. The object of this task is to alternately press the 'a' and 'b' buttons on your keyboard as quickly as possible. You will play for a **<u>total of 10 minutes</u>**.

Every time you successfully press the 'a' and then the 'b' button, you will receive a point. Note that points will only be rewarded when you **<u>alternate button pushes</u>**: just pressing the 'a' or 'b' button without alternating between the two will not result in points. Buttons must be pressed by hand only (key-bindings or automated button-pushing programs/scripts cannot be used) or the task will not be approved.
Feel free to score as many points as you can.

As a bonus, you will be paid an extra 5 cents for every 100 points.

We will divide your 10-minute playing time into 4 rounds.   This means, for example, that if you score 2,000 points, you will receive an extra 1 dollar.

You will have a 10-second break in between each round.   On the next page is an example of how the task will work. Try pressing 'a' and 'b' alternately to score points. We have limited the point total below to a maximum of 50 as this is just practice, but the actual task will not have a limit. Try to reach 50 **as quickly as possible.**

Page Break

JS

Q44 Press 'a' then 'b'...

Points: 0

Page Break

Q45 **You will be paid for your performance.** Please confirm you understand this.

- ○ Yes (1)
- ○ No (2)

Page Break

Q46 Proceed to the next page when you are ready to play the task. Your 10-minute task will begin immediately when the page loads and will be divided into 4 rounds.

Page Break

Button-pushing task

End of Block: Treatment 9 – Piece Rate

Start of Block: Demographics

Q61
Demographics

You are nearly finished. Please answer the remaining demographics questions below.

Page Break

X→

Q62 Did you have any particular strategy when performing the task across these rounds?

- No (0)
- Yes (1)

Display This Question:

If Did you have any particular strategy when performing the task across these rounds? = Yes

Q63 Please briefly describe your strategy.

____

Page Break

Q64 Did you have a points goal or target for each round?

◯ No (1)

◯ Yes (2)

Display This Question:

If Did you have a points goal or target for each round? = Yes

Q65 What was it?

____

Page Break

Q66 How would you agree with the following statements?

I enjoyed this task.

◯ Strongly Disagree (1)

◯ Disagree (2)

◯ Neutral (3)

◯ Agree (4)

◯ Strongly Agree (5)

Page Break

Q67 I felt stress while performing this task.

- ○ Strongly Disagree (1)
- ○ Disagree (2)
- ○ Neutral (3)
- ○ Agree (4)
- ○ Strongly Agree (5)

Page Break

Q68 What is your total household income, including all earners in your household?

▼ Less than $10,000 (1) ... More than $150,00 (12)

Q69 What is the highest level of education you have completed?

▼ Some high school (1) ... Earned graduate or professional degree (7)

End of Block: Demographics