Minds, Models and Markets: How Managerial Cognition Affects Pricing Strategies*

Yi Han David Huffman Yiming Liu

June 6, 2025

Abstract

We show how cognitive constraints systematically shape pricing behavior in a firm with over 20,000 gas stations. Consistent with insights from behavioral models of bounded rationality, station managers with lower cognitive skills have biased mental models. These underestimate competitor sophistication and competitor retaliation for price cuts, and thus place more confidence in the efficacy and profitability of low prices. In the field, managers with lower cognitive skills set lower prices, and engage in more price wars. The resulting lower prices reduce profits and producer surplus, but increase consumer surplus and improve overall market efficiency.

JEL classification: D22, D91, L1

Keywords: Level-k; endogenous depth of reasoning; cognitive hierarchy; mental models; narratives; inductive reasoning; price wars.

*Han: Renmin University of China; yihanecon@ruc.edu.cn. Huffman: University of Pittsburgh; huffmand@pitt.edu. Liu: Humboldt University of Berlin and WZB; yiming.liu@hu-berlin.de.

1 Introduction

Mainstream models of strategic price competition assume that firms are fully rational, and this framework has yielded important insights—such as the ability of firms to sustain high prices in Bertrand-type settings through the threat of retaliatory price cuts. These models abstract from the reality, however, that firm managers are, to varying degrees, boundedly rational.[1] The effects of cognitive constraints are not obvious ex ante: They might help sustain high prices, e.g., by fostering simple strategies that make coordination easier, or they might be neutralized by learning, high stakes, or institutional guardrails within firms. If they work against optimal pricing, two broad mechanisms are possible: firms may pursue optimally high prices but err in implementation—resulting in prices that are too high or too low—or constraints may systematically distort managers' mental models of competition, leading them to adopt fundamentally different strategies, such as favoring price cuts as a way to achieve high profits. If cognitive constraints affect pricing, the implications extend beyond profits to market efficiency, as distortions could change dead weight loss.

In fact, behavioral models posit that bounded rationality can lead to biased mental models of competitor behavior, with implications for pricing. Models such as level-k, cognitive hierarchy, and endogenous depth-of-reasoning (EDR) suggest that individuals engage in limited steps of reasoning about competitor behavior, and then best respond to their predictions (e.g., Nagel, 1995; Costa-Gomes et al., 2001; Camerer et al., 2004; Alaoui and Penta, 2016). This implies a failure of inductive reasoning, as individuals do not generalize from their own behavior to realizing that competitors may think similarly—leading to underestimation of competitor sophistication. The intuition from these models suggests a prediction for real-world pricing: Managers with greater cognitive constraints may charge lower prices, because they find it harder to anticipate competitors' retaliatory price cuts. Laboratory studies link limited inductive reasoning ability to poor performance in strategic games, showing that low scores on the Raven's progressive matrices test predict underestimating competitor sophistication in one-shot "beauty-contest" games and falling into episodes of mutual defection in repeated prisoner's dilemmas (Gill and Prowse, 2016; Fe et al., 2022; Proto et al., 2022).[2] However, field evidence remains scarce.

This paper explores the largely uncharted territory of how bounded rationality of firm managers affects pricing in real-world strategic competition, contributing to the nascent literature on behavioral firms (for a survey see Heidhues and Kőszegi, 2018). We aim to provide field evidence on several sets of open questions. First, do more severe cognitive constraints among firm managers, as captured by lower cognitive skills, lead to biased mental

[1]We use the terms "bounded rationality," "cognitive constraints," and "low cognitive skills," interchangeably.
[2]See also Burnham et al. (2009) and Carpenter et al. (2019) for related evidence.

models of competitors, and if so, does this translate into systematically different mental models of optimal pricing strategies? Second, does this result in differences in actual pricing behavior and engagement in price wars? Third, how do resulting differences in pricing affect firm profits, producer surplus, consumer surplus, and market efficiency and what are the implications of these findings for measuring market power and for competition policy? The data we use to answer these questions come from a collaboration with a company that operates more than 20,000 gas stations.

Section 2 of the paper begins with a conceptual framework for how cognitive constraints may influence managerial pricing decisions. The framework is very general, focusing on a key intuition shared by a range of models of price competition: That the threat of retaliatory price cuts can make high prices optimal. Building on the intuitions of behavioral models, we conceptualize bounded rationality as causing managers to underestimate competitor sophistication, and thus not fully appreciate how competitors will retaliate. This implies greater perceived efficacy of price cuts for increasing both sales and profits, and leads to choosing lower prices. Although our specific application is retail fuel markets, the generality of the framework illustrates that this form of bounded rationality may be relevant for many other types of price competition settings.

The rest of Section 2 gives more details on the market environment, and describes four different data sources we use for our analysis. (1) A survey of approximately 350 district-level senior managers to gather information on the discretion given to station managers on fuel pricing, their views on potential mistakes made by station managers, and the reasons for allowing autonomy despite these mistakes. (2) Multiple survey waves with the 20,000 station managers, achieving roughly 14,000 responses each time, which provide measures of manager cognitive skills including a Raven test, as well as many other manager traits that we use as control variables. The surveys also include measures of mental models of competitor behavior, in the controlled setting of a beauty-contest-type game, survey questions eliciting managers' mental models about their real-world competition, and narratives about optimal strategies for high profits. (3) Four years of monthly panel data on the prices and profits of nearly all of the company's gas stations. (4) For one region, higher-frequency (daily) pricing data for all stations from our partner company, as well as prices of all competitor stations; this allows identifying price wars, and calibrating a simple structural model to quantifying the welfare impacts of boundedly-rational pricing.

Section 3 presents evidence that cognitive skills shape station managers' mental models of competition. We begin with a variant of the beauty-contest game—the money-request game (Arad and Rubinstein, 2012; Fe et al., 2022)—in which managers get whatever amount they request from \$1 to \$6, but with a \$10 bonus for requesting exactly \$1 less than one's opponent. This measure is abstract but places all managers in an identical strategic environment. The modal choice is \$5, suggesting many managers assume unsophisticated opponents, but because so many request this, the optimal choice is \$4 (requests below \$4 lead to lower expected payoffs). The modal manager thus underestimates competitor sophistication. Managers with high cognitive skills, however, are significantly more likely to choose \$4, consistent with ability to correctly anticipate competitor behavior. We next consider two survey questions probing manager's mental models of competitors in their real-world competition. Lower-skill managers are more confident about ability to influence fuel sales, and are more likely to believe that matching competitor prices is optimal – both patterns being consistent with under-appreciating how competitors may retaliate with their own price cuts. A potential concern is that these measures might reflect differences in the real-world strategic environments managers face, e.g., in terms of the numbers of competitors they face, but it turns out market characteristics are quite balanced by manager cognitive skills, and we also find the same relationship of these measures to cognitive skills regardless of variation in market environment.

To examine whether variation in cognitive skills translates into different beliefs about optimal pricing strategies – and does so partly through a mechanism of shaping mental models of competitors – we use a narratives approach (Andre et al., 2023). We ask managers in an open-ended question what they believe contributes most to consistently high fuel profits, and we categorize responses into distinct strategy types, using several robustness checks. A clear pattern emerges: high-skill managers emphasize maintaining high prices (e.g., "Do not blindly engage in price wars"), while low-skill managers stress high volume and low prices (e.g., "Increase sales through price cuts"), consistent with overconfidence in the profitability of cutting prices. Regression analysis confirms that the belief in high-price strategies is significantly associated with cognitive skills, controlling for other manager traits and market conditions. Adding our measures of mental models of competitor behavior to the regression reduces the coefficient on cognitive skills, consistent with these playing a mediating role. The measures of mental models of competitors are also significant predictors: Managers who do not request \$4 in the money request game, or believe they can strongly influence fuel sales, or believe price matching is optimal, are less likely to view high prices as the way to achieve high profits. Taken together, these results are consistent with the conceptual framework's assumption that cognitive constraints bias mental models in the direction of underestimating competitor sophistication and overestimating the efficacy and profitability of price cuts.

Section 4 examines how actual fuel pricing decisions relate to managers' cognitive skills and their mental models. The pricing environment features a government-imposed price ceiling, indexed to world oil prices, which serves as a natural focal point for sustaining high prices. While the two dominant firms—including our partner—often price close to the ceiling, a large number of small, local chains, denoted "independent competitors," are positioned in the market as selling an inferior brand, and typically charge lower prices. We find that lower-skill managers set significantly lower prices on average – an effect comparable in magnitude to adding half an additional independent competitor. These results are robust to controls for other manager traits, station characteristics, location characteristics, and local market structure. Further analysis suggests that the relationship of actual prices to cognitive skills is at least partly mediated by our measures of mental models of competitors. Thus, despite all of the factors that can influence actual pricing decisions in the field, bounded rationality emerges as exerting a systematic influence.

To provide an additional robustness check on causality, and to offer a glimpse into the dynamics of how bounded rationality shapes pricing strategies over time, we conduct an event study. Our main approach is to define treatment events as a station receiving a new manager and estimate how the cognitive skills of the new manager influence the evolution of prices at each treated station, holding everything about the station and market constant. To address time trends we difference with respect to control stations that never have a manager change. We find that after arriving, the pricing strategies of low-skill and high-skill managers diverge over time, with low-skill managers reducing prices and high-skill managers raising prices. New managers' mental models of competitor behavior also matter for pricing in the expected ways. Interestingly, there is no sign of these effects diminishing over the two-year post-period, indicating that the influence of bounded rationality on mental models and pricing is persistent in the face of feedback over substantial time periods.

Section 5 turns to analyzing potential spillovers to market prices more broadly, using data from the region where we observe competitor prices. We show that price cuts by stations of our partner firm trigger price cuts by competitors, and we also show that managers with low cognitive skills tend to respond with deeper price cuts when a competitor cuts price. Both suggest that the pricing strategies of low-skill managers could trigger cycles of sustained and subnormal price competition, i.e., price wars. Creating an indicator for price wars, we find that low-skill managers are involved in about twice as many price wars as high-skill managers, suggesting that at least some price wars may be strategic mistakes.

Section 6 provides additional evidence on mechanisms, drawing on one of our survey waves that more directly elicited manager perceptions of competitor price responses. The survey presented a subsample of managers with real pricing data from a station facing independent company competitors. After our station cut prices, a competitor further lowered its price some days later. We asked managers to predict counterfactual sales of our partner station, but also prices of competitors, if our partner station cut prices on an earlier date. Consistent with our conceptual framework, low-skill managers were significantly less likely to predict a retaliatory price cut. We also see that managers who requested $4 in the money request game were more likely to expect a retaliatory price response. Likewise, expecting a retaliatory price response predicts managers being less confident about ability to influence fuel sales, or about the optimality of price matching. These findings support our interpretation of the mental model measures in terms of relating to underestimation of competitor price responses, and support the explanation of our conceptual framework for why bounded rationality may lead to lower prices.

Section 7 presents evidence on the consequences of the different pricing strategies of low- and high-skill managers for profits and welfare. Profits are negatively correlated with price cuts, and an IV analysis using manager mental models of competitors as instruments for price estimates that a 1 s.d. reduction in cognitive skills reduces monthly profits by about 3% through the channel of lower prices. Another source of evidence is our survey of district managers, which reveals a pervasive concern about a tendency of station managers to charge sub-optimally low prices. Turning to welfare, we estimate demand parameters from a simple, conjectural variation model of differentiated products. We find that a 1 s.d. reduction in cognitive skills translates into as much as a 7% reduction in producer surplus—broadly similar to our reduced-form IV results on profit losses—but generates even larger gains in consumer surplus, thereby reducing dead weight loss by up to 14%. Thus, bounded rationality may improve efficiency by limiting firms' ability to exploit market power. Conversely, another thought experiment is increasing cognitive skills to the maximum observed in our data (a 2 s.d. increase), which approximates the efficiency loss of having full rationality, and could mimic replacing human price setters with algorithms; dead weight loss increases by as much as 28%. We also find that price markups are up to 7% higher for high-skill managers, suggesting that ignoring cognitive skills can bias standard market power measures used in competition policy.

Section 8 uses our survey of district managers to understand why upper management delegates pricing authority to station managers. The main reason cited by district managers is a belief that station managers do have valuable local knowledge about optimal pricing. Consistent with the view that optimal price varies by location, we find substantial heterogeneity in a proxy for station-level demand elasticities. At the same time, data limitations prevent having credible estimates of optimal price at the station level, illustrating why senior management relies on manager local knowledge Another problem with rigid pricing, mentioned by senior management, is that this undermines the ability to threaten price cuts and may thus invite undercutting by competitors. These findings illustrate how bounded rationality can shape pricing even in very large firms, because pricing authority is delegated to local decision makers.

Section 9 provides a concluding discussion. We discuss external validity and implications of our main results. Although our paper is focused on bounded rationality, we briefly discuss some findings on how other manager traits relate to mental models and pricing. We also propose some directions for future research, for example, understanding in more detail why market experience does not seem to lead low-skill managers to change their mental models or pricing strategies.[3] We speculate that the persistence in the face of feedback may indicate mental models that are mis-specified in a way that makes them difficult to falsify, but in future research we hope to explore the specific nature of these alternative models.

[3]Whereas market experience has been shown to eliminate some types of biases, e.g., endowment effects among professional traders (List, 2003).

Our study contributes to literatures on bounded rationality, mental models, and strategic competition. We provide evidence on the importance of insights from the largely theoretical and lab-based literatures on level-k and EDR models for real-world strategic competition.[4] Our results also contribute to the growing empirical literature on misspecified mental models and narratives in economic decisions (e.g., Kendall and Charles, 2022; Andre et al., 2023a; Andre et al., 2023b; Esponda et al., 2024), offering new evidence on how biased mental models vary with cognitive skills, influence real economic decisions and outcomes, and persist in the face of high stakes and feedback.[5] We further add to the literature on behavioral firms (Hortaçsu and Puller, 2008; List and Mason, 2011; Goldfarb and Xiao, 2011 and 2019; DellaVigna and Gentzkow, 2019; Strulov-Shlain, 2023; Tadelis et al., 2023), by directly linking cognitive skills to beliefs, pricing strategies, and performance. Our findings complement lab studies of beliefs and strategic heterogeneity in games (e.g., Dal Bó and Fréchette, 2019; Aoyagi et al., 2024) by showing heterogeneity in real strategic competition and how this relates to cognitive skills. Finally, we contribute to an older tradition of bounded rationality in IO, in which models have assumed that firms use rules of thumb, or face explicit cognitive costs (for a survey see Ellison, 2006), by providing empirical field evidence on how individual-level bounded rationality of managers systematically distorts mental models of competitors and pricing.

Our findings speak to more mainstream literatures on strategic competition. They add nuance to the view that price wars signal collusion (e.g., Green and Porter, 1984; Slade, 1992), by showing that some may instead result from cognitive mistakes. By identifying bounded rationality as a factor influencing price markups, we contribute to the literature on measuring market power (see Berry et al., 2019). We also add to evidence that price ceilings can serve as focal points (e.g., Knittel and Stango, 2003), while showing that this depends on managers' cognitive skills. Our results complement work on competition in retail gas markets (e.g., Hastings, 2004; Noel, 2007; Barron et al., 2008; Houde, 2012; Luco, 2019). For instance, Assad et al. (2023) find that algorithmic pricing raises fuel prices in Germany; our findings suggest a reason why—that human pricing may leave market power partially unexploited.[6]

Our findings contribute to an economics literature showing that managers affect worker and firm performance (e.g., Ichniowski et al., 1997; Bloom and Van Reenen, 2007; Bloom et al., 2013, 2019; Bandiera et al., 2020; Hoffman et al., 2021; Fenizia, 2022; Adhvaryu et al., 2023; Metcalfe et al., 2023; Minni et al., 2023). While most of this work focuses on

[4]Our findings are also consistent with lab evidence that implementing certain strategies may be cognitively demanding (Oprea, 2020; Proto et al., 2020).

[5]Our results also complement recent research in behavioral finance and macroeconomics on the importance of "partial equilibrium thinking" among professionals and consumers (e.g., Bastianello and Fontanier, 2024).

[6]By shedding new light on the properties of human price setting, our findings also complement a growing literature that has used simulations and laboratory studies to investigate the properties of algorithmic pricing (e.g., Calvano et al., 2020; Pai and Hansen, 2020; Asker et al., 2022 and 2024; Fish et al., 2024; Arunachaleswaran et al., 2024).

supervisory roles, we have an unusually large sample of managers making outward-looking strategic decisions and study the role of cognitive skills in shaping mental models, pricing, and profits. Prior work links overconfidence to CEO investment choices (Malmendier and Tate, 2015) and shows that such overconfidence can persist despite feedback (Huffman et al., 2023). Our results suggest that bounded rationality of managers can foster underestimation of competitor sophistication, and thus overconfidence about ability to profitably compete using price cuts.

Our paper also complements literatures in labor economics and psychology on how cognitive skills predict wages and job performance (e.g., Boissiere et al., 1985; Cawley et al., 2001; Schmidt and Hunter, 2004; Heckman et al., 2006). We provide some of the first evidence on how cognitive constraints relate to a key decision in strategic management roles, price-setting, for which the impact is not obvious ex ante. Furthermore, we show that the underlying mechanism is not just noise, but at least partly due to biased mental models, which points to a systematic tendency to set lower prices, and implies a *positive* welfare consequence of low cognitive skills that mitigates market failures arising from market power.[7]

2 Market setting and Data

2.1 Conceptual framework

This section outlines a general conceptual framework for how bounded rationality can influence pricing, focusing on a key comparative static common to many models. The generality highlights that the mechanism we propose may apply broadly, beyond the retail fuel markets we study.

We assume a market with multiple firms, each of which has a manager who determines pricing policy and is motivated by profits.[8] Drawing on a key insight from mainstream theories of strategic competition, we assume that the ability to threaten punishments, in the form of retaliatory price cuts, is a source of market power. We denote competitor retaliation by $\frac{\Delta p_{-i}}{\Delta p_i} > 0$, where p_{-i} is competitor price and p_i is own price, and the competitor price response may or may not be a continuous function of the change in own price. As is well known, if competitor retaliation has sufficiently severe negative consequences for own profits, and the firm values the future sufficiently, the optimal price can be higher than it would be in the absence of anticipated punishments, potentially as high as the monopoly price.

A key comparative static, which can be generated by a range of more specific models, is

[7]Our findings show a novel way in which bounded rationality can improve market efficiency, complementing previous evidence that bounded rationality can be beneficial, e.g., by reducing dead weight loss from taxes (Chetty et al., 2009), or limiting surplus-reducing exploitation of loopholes in workplace incentive schemes (Abeler et al., 2025).

[8]Managers may not be motivated solely by profits, but if profits enter positively in their objective function, this provides a reason why they can be influenced by the threat of profit-reducing competitor price cuts.

that a weaker perceived competitor price response increases the perceived profitability of price cuts and makes it more likely that a low price is chosen, all else equal. One example is a standard repeated Bertrand competition model, combined with the assumption of an equilibrium in grim trigger strategies (e.g., Deneckere, 1983). If a manager does not anticipate, in the current period, the grim trigger response to a price cut, or does not fully anticipate the depth of the price response, this will lead to choosing a lower price in the current period.[9] Another example, which abstracts away from the complication of a repeated game framework, and has a continuous response of competitor price to own price, is a classic conjectural variation framework in the spirit of Porter (1983) and Slade (1984). In this one-shot framework, which can be thought of as a reduced form for a repeated game under certain conditions, the firm incorporates a perceived competitor price "response" into the profit maximization problem, $\theta_i = \frac{dp_{-i}}{dp_i}$ (for conditions see Corts, 1999). It has the comparative static that a lower perceived θ_i (weaker competitor price response) leads to greater perceived profitability of price cuts and the choice of a lower price.[10]

We conceptualize bounded rationality of the firm manager as influencing pricing decisions because of how it shapes mental models of competitor behavior. Behavioral models such as level-k reasoning and EDR, and empirical evidence from laboratory studies, suggest that individuals with a trait of lower cognitive skills find it harder to think deeply about how competitors respond to the game, and take competitor behavior as given rather than accounting for the possibility that competitors may think and act like them. We assume that this means that boundedly rational managers perceive a weaker competitor price response, $\alpha_i \frac{\Delta p_{-i}}{\Delta p_i}$, with $0 \leq \alpha_i < 1$, and α_i decreasing for lower levels of cognitive skills. This in turn implies a tendency to choose lower prices.[11]

The empirical relationship that we seek to test is whether lower cognitive skills lead to lower chosen prices. We will also test for elements of the hypothesized mechanisms linking

---

[9] In the one-shot Bertrand model, with symmetric firms, firm $i \in \{1,2\}$ faces demand $q_i = a + bp_i + cp_{-i} + g(z)$, where q_i is quantity sold, and demand depends on own price, competitor price, and demand shock $g(z)$. Firm i chooses price to maximize $\pi_i = (p_i - MC)q_i$. The (interior) Nash equilibrium prices are $p_i^* = \frac{b \cdot MC - a - g(z)}{2b+c}$. In the repeated game with discount factor δ, firms may sustain collusive prices $p^C > p^*$ using grim-trigger strategies. The incentive compatibility constraint is $\frac{\pi^C}{1-\delta} \geq \pi^D + \delta \frac{\pi^*}{1-\delta}$, where π^C, π^D, and π^* denote collusive, deviation, and Nash profits, respectively. A fully rational firm manager anticipates $\Delta p_{-i} = p^* - p^C$ in response to $p_i < p^C$, and profit π^* in all future periods, whereas a boundedly rational manager anticipates $\alpha \Delta p_{-i}$ and $\pi > \pi^*$. Thus, the incentive compatibility constraint is less likely to hold for a boundedly rational manager and they are more likely to choose a price lower than p^C.

[10] Residual demand for firm i is again given by $q_i = a + bp_i + cp_{-i} + g(z)$. Firm i chooses price to maximize $\pi_i = (p_i - MC)q_i$. The first-order condition is $p_i \frac{dq_i}{dp_i} + q_i = MC \cdot \frac{dq_i}{dp_i}$, where $\frac{dq_i}{dp_i} = b + c\theta_i$ and $\theta_i = \frac{dp_{-i}}{dp_i}$ reflects station i's conjecture about competitor responses. Solving yields optimal price $p_i^* = \frac{MC(b-c\theta_i)-a-c\cdot p_{-i}-g(z)}{2b-c\theta_i}$. The derivative is $\frac{dp_i^*}{d\theta_i} = \frac{c[a+b\cdot MC+c\cdot p_{-i}+g(z)]}{(2b-c\theta_i)^2}$. For substitute products, $c > 0$, and the expression $a+b\cdot MC+c\cdot p_{-i}+g(z)$ represents the quantity demanded when price equals marginal cost, which must be positive for the market to exist. Thus, $\frac{dp_i^*}{d\theta_i} > 0$.

[11] Recent work by Aoyagi et al. (2024) proposes a model of level-k players in an indefinitely repeated PD game, where bounded rationality fosters defection by leading to pessimistic beliefs about the possibility of sustaining cooperation. This is another example of how incorporating misunderstanding of competitor responses can lead to a prediction of more aggressive competition.

bounded rationality to lower prices: underestimation of competitor sophistication, leading to underestimation of their price responses, which in turn increases perceived profitability of price cuts.

2.2 Details on the market setting and manager descriptives

Our partner company operates more than 20,000 gas stations across a country. Stations typically sell both gas and diesel fuel, and essentially always have a convenience store. In this country and company, the bulk of station profits come from fuel sales rather than the convenience store; on average in our dataset, fuel profits make up 71% of total station profits. The stations are primarily company-owned, rather than franchises. Each station has a station manager, who has substantial influence over station operations, including pricing decisions. Station operation is also governed, however, by the policies of more senior, district-level managers. There are about 350 districts, each with a district-level manager who sets policies about precisely what type and degree of discretion is given to station managers operating in their district.

The market is relatively concentrated, and also has important types of brand differentiation. Our partner company is one of two major brands in the market for retail gasoline, with stations all across the country and each claiming about one third of the market. The rest of the market share goes to hundreds of smaller, more local chains, which we refer to as "independent" competitors. One key difference between the major brands and independent brands is that the former produce their own fuel, while the latter must buy fuel products on the market. This difference in vertical integration becomes an important feature in estimating demand elasticity through cost shifters in Section 7. Another key difference is that the larger companies position themselves as offering a premium fuel product, and for this reason typically charge more than the independent companies for the same grade of gas or diesel.

An important feature of the pricing environment is that there is a government-imposed price ceiling for fuel products, indexed to the world price of oil. The price ceiling arguably serves as a natural focal point for coordinating pricing, and indeed, the two large companies have a policy of generally pricing near the price ceiling. Independent stations also frequently exhibit pricing that tracks movements in the price ceiling, albeit with a substantial discount.

Table 1 shows descriptive statistics for station managers and their stations. The median age of a station manager is 39, and about 70 percent of managers are male. The modal level of education is a junior-college degree. Managers stay in their jobs for a substantial period of time, with median experience at the company being 7 years. Managers do switch gas stations periodically, with median tenure at a gas station of about 2 years. The median number of employees is 5, so managers have some people management duties, but not for very large groups of workers. The median number of competitors within the local market is 3.

Table 1: Descriptive Statistics: Managers and Stations

Manager descriptives		Station descriptives	
Median age	39	Median number of employees	5
Female	34%	Median number of competitors	2
Education level:		Median market share	30%
High school	26%		
Junior college	45%		
College or above	28%		
Median experience (years)	7		
Median tenure at current station (years)	2		

Notes: This table reports descriptive statistics based on the first survey wave. "Experience (years)" represents the total number of years as a station manager, while "tenure at current station (years)" represents the number of years as the station manager at their current assignment. "Competitors in local market" generally refers to stations from other companies within 2.5 km, though the definition also takes into account road configurations, as determined by the partner company. "Market share (of fuel sales)" is a station's share of total fuel sales in its local market as reported by the station manager.

The company's definition of the local market starts with a 2.5 km radius around the station as a guideline, but then calls for adjusting the assessed number and type of competitors based on considerations like commuting routes, etc. The median market share for a station from our partner company, in terms of fuel sales in the local market, is about 30 percent.

2.3 Datasets

We obtained data through a collaboration with the partner company's research arm.[12] This includes performance data (e.g., profits and prices), but also survey data collected via the company's internal survey infrastructure. Because participation in internal surveys is expected—but these are administered by a unit independent of senior management—response rates are high and confidentiality is credible. Data from the human resources department of the company were not available. For this reason, our surveys were designed to collect variables that describe some aspects of the work environment that would normally be collected by human resources, such as manager work history. We were not allowed to collect certain variables, however, such as manager earnings or work hours.

Our analysis is based on four types of datasets, the first of which is from a survey conducted with senior, district-level managers. We have responses from 353, close to a 100 percent response rate. One purpose of this survey was to collect systematic information about the amount of discretion given to station managers over fuel pricing. Another was to elicit senior manager views on potential mistakes by station managers when it comes to pricing.

A second type of dataset comes from surveys conducted with the station managers. We have conducted three survey waves between 2022 and 2024, each time sending the survey

[12] IRB approval for this study was granted by Renmin University of China.

to all 20,000 station managers. The response rate has been consistently around 70 percent, yielding approximately 14,000 responses per wave. Our first survey, conducted in 2022, collected measures of a wide array of manager traits, as well as measures of managers' mental models of competition. The second survey, which was completed in 2023, collected measures of the same traits again, but also included additional measures of mental models, most notably a measure of managers' "narratives" about the causes of high profits. The third wave, completed at the end of 2024, included a measure, implemented for a random subsample of managers, which was designed to directly test whether managers recognize a causal link between their own price cuts and competitors' price responses.

A third dataset is monthly performance data on the company's gas stations, for the period 2019 to 2022. These are panel data for each station, recording key outcomes such as fuel and nonfuel profits, sales volume in gallons, and average monthly prices charged for gas and diesel products. We have access to the monthly performance data in 26 of the 31 regions the company operates.[13] The total dataset has about 16,000 gas stations. The research arm matched performance data to responses from our first survey wave using an internal company identifier, yielding linked data for roughly 10,000 stations (due to the 70 percent survey response rate). The second and third survey waves have not been similarly matched by the company, but we have matched responses across survey waves ourselves using other identifiers included in the surveys. This allows linking second- and third-wave responses to performance data, but only for managers who also responded to the first wave and were matched to performance data by the company.

A fourth dataset includes daily price and sales data for one region, including prices and characteristics of all competitors. The region has roughly 900 stations from our partner company. The data form a daily panel for each station, from 2018 to 2021. These data allow identifying price wars because of information on competitor prices. The high-frequency data also facilitate estimating demand parameters for a structural model that we use for welfare analysis.

Table 5 in the Appendix summarizes the data structure by showing which types of measures are included in our different survey waves with managers, and in the administrative dataset on station performance.

2.4 Managerial discretion over fuel prices

Our survey of district managers shows that station managers have substantial influence over fuel prices, although the degree of autonomy varies across districts. In 48% of the districts, the senior manager reports that station managers can directly change listed fuel prices without needing to submit a proposal (either within a pre-specified range or without restriction).

[13] Of the remaining five regions, some only record data on a quarterly basis, and others were not made available due to idiosyncratic bureaucratic factors.

In the remaining districts, managers must make proposals to change listed prices. The survey of district-level managers indicates an average approval rate of about 34 percent. Thus, even when proposals are required, managers can still influence pricing. Overall, station managers have a clear role in influencing listed fuel prices, opening up the possibility that bounded rationality may matter for pricing and performance. On the other hand, manager autonomy is not without limits, and as we discuss in Section 9, the existing restrictions on autonomy partly reflect senior managers trading off benefits of manager local knowledge against a perceived tendency for station managers to set prices too low.

2.5 Manager incentives

Managers receive a base salary and substantial performance-based pay, which accounts for roughly 50% of total compensation. Bonuses are tied to three KPIs: fuel profits, nonfuel profits, and sales volume. The goal of the company is to incentivize profits, and as such, the sales KPI is mainly intended to encourage good customer service, which can increase sales and profits for a given price. Performance is assessed relative to targets, with weighted contributions to bonus pay. We do not observe manager earnings, or performance targets, as these are internal HR data, and thus cannot recover realized compensation.

Knowing the structure of the scheme, however, makes clear that it meets the requirements of our conceptual framework: managers have a motive to care about profit. This means that our conceptualization of bounded rationality has bite. Specifically, managers who have high cognitive skills will be more likely to perceive that price cuts reduce fuel profits, since they expect competitor retaliation to dampen gains on the margin of increasing sales. Given that their bonus depends partly on fuel profits, they have a reason to internalize this tradeoff and potentially limit price cuts. Managers with lower cognitive skills, by contrast, are assumed not to fully appreciate competitor retaliation, which may weaken or even eliminate the tradeoff they see from cutting prices; they believe price cuts will strongly increase sales, potentially enough to even increase fuel profits.[14]

2.6 Measures of manager traits and construction of cognitive and noncognitive skills factors

Our first and second survey waves collected rich data on manager characteristics via online surveys administered by the company's research department, with roughly 14,000 respondents each time. Table 2 summarizes all traits, beginning with cognitive ability. This section

[14]Managers of both types may choose a price that is lower than the profit-maximizing price, because they realize that lowering price is another way, in addition to good customer service, to increase their bonus from sales. Thus, the incentives may not fully align manager goals with the company goal of maximizing profits. Managers with low cognitive skills are predicted, however, to choose a lower price than managers with high cognitive skills. We do not make any claim that the firm's incentive scheme is optimal.

describes the traits we measure, and then our approach of using factor analysis to reduce dimensionality and uncover latent traits underlying collections of measures, yielding our cognitive and noncognitive skill measures.

Table 2: Measures of manager traits

Cognitive ability	IQ test involving 9 progressive Raven's matrices (+)
Numeracy	Question about understanding probabilities (+)
Economic preferences	Risk tol. (+), patience (+), altruism (+), pos. rec. (+), neg rec. (-)
Ambiguity aversion	Prefer urn with known vs. unknown distributions (-)
Personality type	Big5: Consc. (+), agree. (+), extra. (+), open. (+), neur. (-)
Locus of control	Inventory from psychology (+)
Competitiveness	On a scale from "not at all" to "very" (+)
Confidence	On a scale from "not at all" to "very" (+)
Procrastination	Agreement on a scale about tendency to procrastinate (-)
Liking for authority	On a scale from "not at all" to "very much" (-)
Self control	Inventory from psychology (+)
Emotional intelligence	8 item Reading the Mind in the Eyes Test (+)
Gender	Female indicator
Age	In years
Experience	In months as station manager

Notes: Cognitive skills are measured by the first factor of first two measures in the table (colored in red). Noncognitive skills are measured by the first factor of the measures from economic preferences to emotional intelligence (colored in blue). The signs of factor loadings for the first survey wave are shown in parentheses.

Cognitive ability measures: Our main measure of cognitive ability is a 9-item Raven test. While the full version includes 60 questions, the abbreviated 9-item version has been shown to be a reliable proxy (Bilker et al., 2012). The Raven test is widely viewed as the leading measure of inductive reasoning ability, i.e., ability to generalize from an observed pattern to a general rule, which is a key element of fluid intelligence (Carpenter et al., 1990).[15] Each item presents a 3×3 matrix of visual patterns that evolve across rows or columns according to a set of underlying rules; respondents must infer the rules that generate the observed variation (see Appendix Figure 12 for an example item). We selected the Raven test partly because Laboratory studies show that it predicts performance across a range of different strategic games where success requires inductive reasoning about competitors. This is consistent with the test measuring a broadly relevant trait of bounded rationality. The Raven test is also conceptually related to the notion of bounded rationality in level-k and EDR-type models, where limits on reasoning have an aspect of failure of induction. For our particular application – strategic price competition in fuel markets – it seems plausible that inductive reasoning ability is a crucial determinant in station managers' abilities to form accurate mental models of their competitors. Besides capturing the depth of reasoning about how competitors respond to a given set of incentives, it may capture managers' ability to

[15]Carpenter et al. (1990) describe this form of intelligence as "the ability to reason and solve problems involving new information, without relying extensively on an explicit base of declarative knowledge derived from either schooling or previous experience."

generalize from the data they observe in their real markets to form accurate mental models of general rules governing competitor behavior.

A second measure of cognitive ability is a question about the probability of a flipped coin landing heads, serving as a proxy for numeracy, an aspect of crystallized intelligence.[16] Responses turn out to be modestly correlated with Raven test scores ($\rho = 0.14$; $p < 0.001$), and both load on the same main factor in our factor analysis, described below.

Other manager characteristics: In selecting other traits to measure, we sought to cast a wide net, selecting a wide range of measures that are viewed by economists as capturing important drivers of economic decision making and by psychologists as key facets of human nature. Six aspects of preferences—risk, time, and various social preferences—were measured using the survey module from the Global Preference Survey (Falk et al., 2018). These survey measures were developed based on their ability to predict choices in incentivized experiments measuring the corresponding preferences. Personality type was measured using a Big Five personality inventory. Other items captured beliefs and other types of preferences, including locus of control (German SocioEconomic Panel Study, SOEPv28 English version), self-reported confidence, taste for competition, taste for authority, self-reported self-control (Tangney et al., 2004), procrastination, and ambiguity aversion (Dimmock et al., 2016). The measure of emotional intelligence was an eight-question test, showing respondents photographs of a person's eyes and asking them to identify the person's facial expression (Baron-Cohen et al., 2001).

Factor analysis: To reduce dimensionality and capture latent traits underlying groups of trait measures, we performed factor analysis. More details are provided in Appendix A.3. An initial analysis pooling all traits showed a distinction between the cognitive ability measures, which load mainly on one factor, and the other manager traits, which load mainly on other factors. This structure is quite consistent across both the first and second survey waves (Tables A.1 and A.2). Following the terminology of Heckman et al. (2006), we refer to this distinction as corresponding to "cognitive skills" and "noncognitive skills." Performing factor analysis on the cognitive ability measures, we find that both the Raven test and numeracy items load onto a single factor with eigenvalue greater than 1. This is robust across survey waves, with loadings quite similar across waves (Table A.3). We use the corresponding first factor as our cognitive skills measure for each wave. Similarly, noncognitive traits yield a single dominant factor in each survey wave, with a high eigenvalue in both waves and similar factor loadings across waves (Table A.4). This factor loads positively on traits like conscientiousness, agreeableness, locus of control, confidence, and patience, and negatively on traits such as neuroticism, taste for authority, and procrastination. A second noncognitive factor emerges in each wave, loading on traits such as risk taking, openness, and taste for

[16]Crystallized intelligence is conceptualized as distinct from fluid intelligence, as it draws on accumulated knowledge (Cattell, 1987).

authority, but its eigenvalue is near or below 1. We therefore focus on the first noncognitive factor from each wave as our corresponding noncognitive skills measure.

Robustness of measurement approach: We conduct extensive robustness checks to assess whether the use of factors affects our results. We re-run all of our main analyses and find similar results with each of the following adjustments: (1) using Raven scores alone as the measure of cognitive skills, treating the numeracy measure as a separate control variable; (2) controlling for all other manager traits individually rather than as a noncognitive factor; (3) including the second noncognitive factor as another control; or (4) incorporating emotional intelligence into the cognitive factor (results of all of these checks are available upon request). Measurement error in cognitive skills, our key explanatory variable, works against finding significant effects, but a potential concern is that measurement error in our noncognitive skills measure might bias the coefficient on cognitive skills upwards if we use these as controls in regression analysis (Gillen et al., 2019). Using our two waves of trait data to correct for measurement error, however, we find no evidence that this inflates coefficients on cognitive skills.[17]

3 Cognitive skills and mental models

This section examines whether lower cognitive skills bias mental models of competitors, leading to underestimation of sophistication and price responsiveness, and whether this in turn supports the belief that low prices are optimal, consistent with our conceptual framework. To assess mental models, we avoided direct questions about price cuts causing competitor price reactions, which could trigger respondents to think of such reactions even though they typically neglect these in their daily decision making. Instead, we used four more indirect approaches. In the first two survey waves we included: (i) an abstract strategic game, where choices can reveal anticipation of competitor behavior; (ii) survey questions measuring beliefs about real fuel market competition, where underestimating competitor reactions implies predictable patterns in beliefs; and (iii) an open-ended question about fuel profit drivers, which allows spontaneous responses about benefits of high or low prices. In our third survey wave, we added a fourth approach for a subsample of managers: presenting real station data and asking managers to predict competitor pricing following a price cut (we discuss this in Section 6).

[17]Trait measures collected across waves allow us to assess attenuation. Correlations with cognitive skills are attenuated by about 35 percent, and with noncognitive skills by 17 percent, indicating that the estimates we report in our main analysis are lower bounds. Appendix A.4 provides details on robustness to measurement error in controls using the Obviously Related Instrumental Variables (ORIV) method (Gillen et al., 2019; see also Stango and Zinman, 2020).

3.1 Money request game

In our surveys, managers played a hypothetical version of the money request game (Arad and Rubinstein, 2012; Fe et al., 2022): "Suppose you are matched with another station manager to play a game. Your opponent and you are going to ask for an amount of money from a referee for the game. The amount must be between \$1 and \$6. You will get the amount of money you ask for. However, you will get \$10 more if you ask for exactly \$1 less than your opponent. How much money do you ask for?" In level-k and EDR models, lower requests in this game reflect deeper reasoning.[18] More broadly, the game tests strategic understanding and anticipation of others' behavior. A key advantage of the game is isolating strategic sophistication in a controlled environment, though it is abstract and distinct from the competitive settings managers actually face.

Figure 1: Behavior in the money request game and cognitive skills

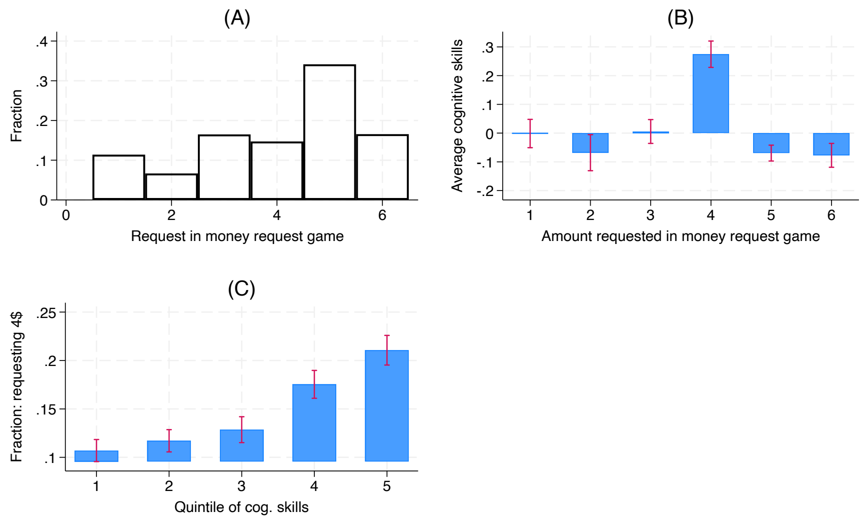

Notes: Panel (A) shows the distribution of requests in the money request game from our first survey wave. Panel (B) shows average cognitive skills for the group of managers making each of the possible requests. Panel (C) shows the share of managers requesting \$4 by quintiles of cognitive skills (5 is the highest). Error bars indicate 95% confidence intervals.

Figure 1 shows results from the money request game. Panel (A) presents the distribution of requests for the roughly 13,500 managers who participated. The modal request is \$5, consistent with anticipating that many others will act non-strategically and request \$6. Because so many managers think this way and request \$5, however, the expected payoff is maximized by requesting \$4 (requests below \$4 lead to lower expected payoffs). This means that the modal manager makes a mistake of underestimating how many others choose \$5 as they do. One might think that requesting \$4 is a matter of luck, but Panel (B) of Fig-

[18] The game has no pure strategy Nash equilibrium. The mixed strategy equilibrium involves a mixing distribution of 0, 0, 0.4, 0.3, 0.2, and 0.1 for requests 1 to 6 (Fe et al., 2022), which does not resemble the observed distribution for station managers.

ure 1 shows that average cognitive skills are significantly higher among managers choosing \$4 compared to those making any other request, and Panel (C) further shows that the frequency of choosing \$4 is monotonically increasing in quintiles of manager cognitive skills. Thus, managers with higher cognitive skills are more able to correctly anticipate the degree of competitor sophistication and respond accordingly.

Regression analysis shows that the probability of requesting \$4 is significantly increasing in cognitive skills even when controlling for other manager traits (noncognitive skills, experience, gender, and age) as well as market, location and station characteristics (see the first coefficient in Figure 2 from a linear probability model). Taken together, our results indicate that cognitive skills are important for station managers to have an accurate understanding of the strategic behavior of other station managers, in an abstract but tightly controlled setting.

Figure 2: Mental models of competition and cog. skills: Controlling for other traits and market conditions

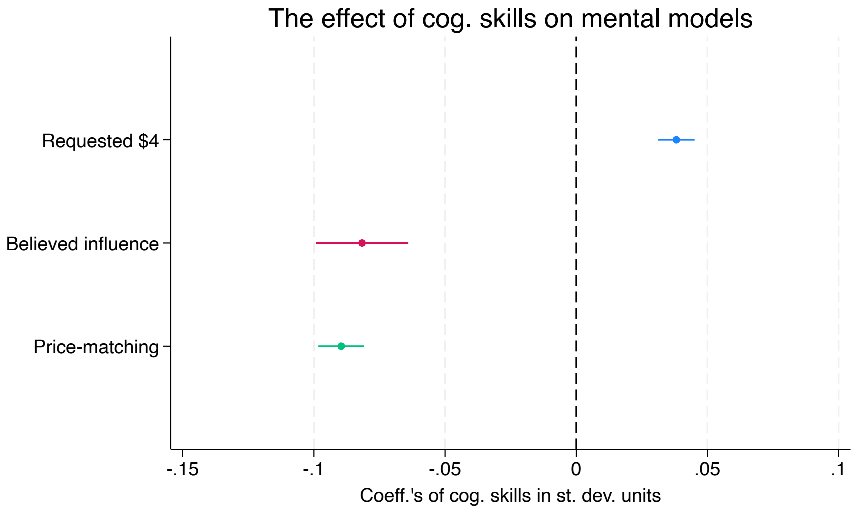

Notes: This figure reports coefficients from OLS regressions on three mental model measures, with 95% confidence intervals. The three dependent variables are: requesting \$4 in the money request game, believed influence over fuel sales, and believed optimality of price-matching strategy. Coefficients are for cognitive skills as an explanatory variable. Each regression controls for other traits of managers (noncognitive skills, age, gender, experience), and station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators, district fixed effects). Appendix Table B.2 provides the underlying regression estimates.

Robustness: As detailed in Appendix B.1, the findings on the money request game are strikingly robust across surveys waves. In each of our second and third survey waves, \$5 was the modal request, \$4 was the expected-payoff-maximizing choice, and those requesting it consistently displayed the highest cognitive skills (see Appendix B.1).

3.2 Beliefs about influencing fuel sales and optimality of price matching

We also included less abstract measures of managers' strategic thinking in real market environments. One set of questions asked: "Compared to objective factors such as the location of the gas station and the number of competitors, in your opinion, to what extent can managers influence the sales of fuel products [overall performance; convenience store profits] of a station?"[19] Based on the conceptual framework, we hypothesize that lower cognitive skills may lead to *greater* confidence in influencing fuel sales, because of not fully anticipating how the effects of price cuts on sales will be dampened by competitor retaliation. Although the inference that confidence comes from belief in the efficacy of price cuts is indirect, we will examine this more directly below. Another question asked managers: "Last month, you believed that the optimal price would be to match a competitor's price."[20] In the terminology of our partner firm, price-matching refers to matching a competitor's price when they are pricing below the ceiling, otherwise the pricing policy is described as setting price at the ceiling.[21] A belief that price-matching is optimal is therefore an indicator of willingness to cut prices, and if this reflects failing to anticipate competitor retaliation, one would expect the frequency to be higher for managers with lower cognitive skills.

Figure 3 shows how these measures relate to cognitive skills. Panel (A) shows that managers with lower cognitive skills are indeed more confident in the ability of managers to influence fuel sales, whereas cognitive skills are more weakly related to beliefs about influencing overall performance, and largely unrelated to beliefs about influencing convenience store profits. This is consistent with a mechanism specifically related to underestimating competitor sophistication in fuel price competition. Panel (B) shows that favoring a price-matching strategy decreases monotonically with cognitive skill quintile. Regression analysis (see second and third coefficients in Figure 2) shows that these relationships are statistically significant and robust to controls for other manager traits as well as station and market characteristics.

Robustness: A potential concern is that these results might not reflect a causal impact of bounded rationality on beliefs, but rather managers with different cognitive skills facing

[19]We asked about ability to influence overall performance and convenience store profits to check whether, as expected, cognitive skills would matter mainly for confidence about fuel sales.

[20]The wording includes a reference to the previous month because, to reduce salience of the measure, we embedded it in an inventory of management practices based on McKenzie and Woodruff (2017) with slight wording changes to fit the gas station environment, and their inventory asks about a management practice being employed in a specific time frame. While our analysis focuses on price-matching behavior, we also examined other price-related items in the original inventory, specifically "Visited competitor's stations to see prices" and "Used a customized offer to attract new customers." We find that managers with high cognitive skills are less likely to report visiting competitor stations to see their prices and are more likely to report attracting new customers through customized offers.

[21]Consistent with this interpretation, we show below that managers who believe price-matching is optimal are less likely to believe that high prices are beneficial for profits.

Figure 3: Beliefs about influencing performance and price-matching, by cognitive skills

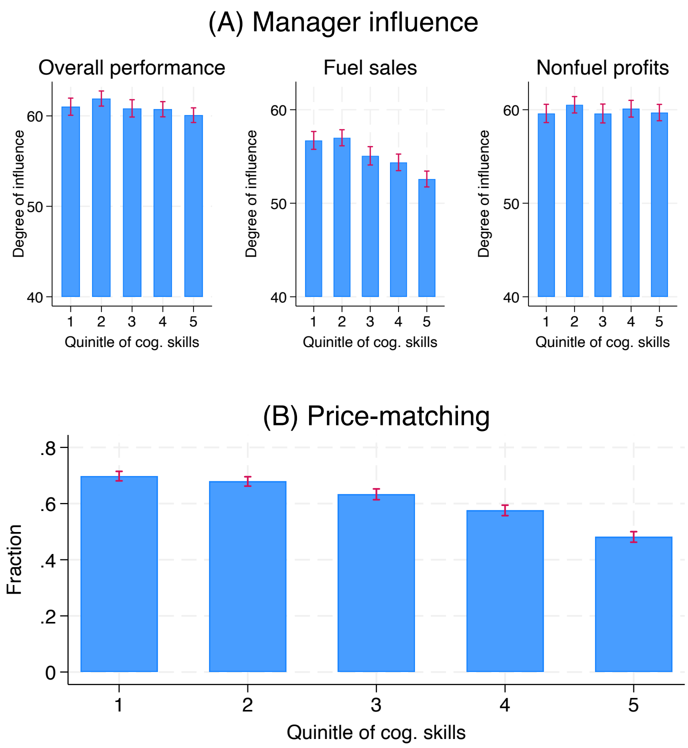

Notes: Panel (A) shows beliefs about manager influence measured on a scale from 0 (only external factors matter) to 100 (only manager matters). The three graphs show perceived manager influence on overall performance, fuel sales, and nonfuel profits respectively, by cognitive skill quintile. Quintile 5 is the highest. Panel (B) shows the share of managers believing that matching competitor prices is optimal. Error bars indicate 95% confidence intervals.

systematically different market conditions, which happen to lead to learning these different viewpoints. However, Appendix Table B.1 shows that market conditions are actually very similar for high- and low-skill managers; differences in location type or number of competitors are economically small. Moreover, as was shown in Figure 2, results are robust to including a rich set of controls for market structure, location type, and other manager traits. Appendix B.2 further shows that the relationship between cognitive skills and these mental model measures also holds within different kinds of market conditions and locations.

3.3 Narratives about determinants of fuel profits

To examine whether cognitive skills shape higher-level mental models of profit determinants—potentially by influencing mental models of competitor behavior—we used a "narratives" approach following Andre et al. (2023). In our second survey wave, we asked: "Some managers consistently have higher fuel profits than other managers. What do you think are the most important practices that enable them to achieve this?" This open-ended question elicits managers' conceptualizations of profit drivers without imposing structure.

As detailed in Appendix B.3, most responses could be mapped onto the profit formula (profit = profit-margin × sales), with managers emphasizing either margins (e.g., maintaining high prices) or sales volume (e.g., boosting volume through low prices). Some responses, such as effort or ability, did not fit cleanly into either. We developed a classification rubric with 15 categories (see Appendix H). Two RAs, blind to our hypotheses, used the rubric and independently classified the more than 15,000 responses, with a 75 percent agreement rate between RAs; in our best attempt to reduce measurement error, we report results based on researchers reconciling disagreements, but as described below, we check robustness to alternative approaches. In the classification exercise, most managers (80%) mentioned only one cause, some (17%) mentioned two, and very few mentioned more.

Figure 4 presents the frequencies of mentioning different types of causes and their relationship to cognitive skills. Panel (A) shows 25% of managers cite profit margin causes, 45% mention sales volume, and 30% cite other causes or "don't know." Panel (B) ranks causes by the cognitive skills of those who mention them, revealing a pattern that those citing high profit margin causes (high price is by far the most frequent of these) generally tend to have have higher skills, while those focusing on sales volume or low prices tend to have lower skills. Panel (C), which excludes infrequent causes, sharpens the picture: cognitive skills are highest among those citing high price, followed by sales volume and low price, and lowest among those citing ability, effort, or "don't know." The differences in cognitive skills, comparing managers mentioning high price to those mentioning low price, or sales volume without elaboration, or customer development are statistically significant (Wilcoxon tests; $p < 0.001$).[22] Regression results (Figure 5) show that the probability of mentioning high price is significantly increasing in cognitive skills, controlling for other manager traits and market and station characteristics. Thus, we see different views emerging from the data about how to achieve the same goal of high fuel profits, one focused on high price and the other focused on sales volume and low prices. Consistent with the conceptual framework, boundedly rational managers are less likely to think that high prices lead to high profits. A caveat with the narratives measure is that there are likely other reasons why managers mention high price besides thinking about competitor behavior. For this reason, we turn to investigating whether mentioning high price is related to mental models of competitor behavior.

We use regression analysis to examine whether mental models of competitor sophistication help explain views on the determinants of profit, and whether these are mediators in the relationship between cognitive skills and profit narratives. We find that each of the mental model measures is a significant predictor of the high price narrative, controlling for cognitive skills (see Columns (2) to (4) in Table B.3). Managers who had better insights into competitor behavior in the money request game (requesting \$4) are more likely to fa-

[22]For these tests we exclude managers who mention multiple causes to ensure independence of observations.

Figure 4: Narrative measure of mental models for high fuel profits

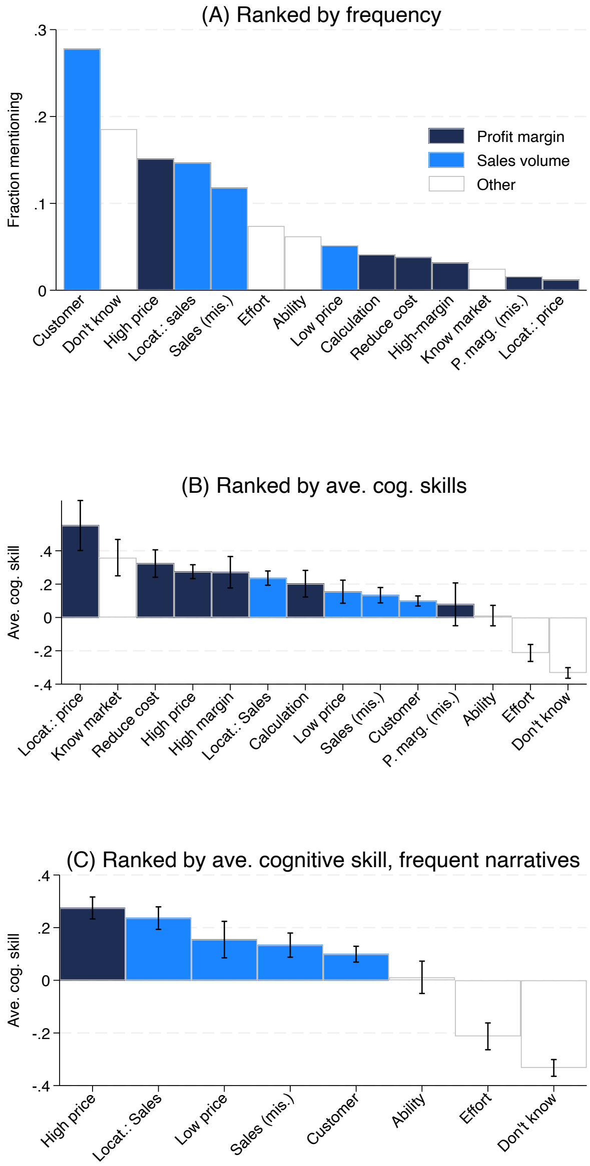

Notes: Panel (A) shows the frequencies of managers mentioning different categories of causes of high fuel profits. The bars are color-coded according to whether they fall into the profit margin, sales volume, or other categories. The narratives include: *Customer* (good customer service); *Don't know* (no clear explanation offered); *High price* (maintaining a high price); *Location: Sales* (the location is favorable for high volume); *Sales (mis.)* (mentioning sales volume but without further explanation); *Effort* (manager's hard work); *Ability* (manager's high abilities); *Low price* (generating high volume through low prices); *Calculation* (calculating benefits and costs before deciding what to do); *Reduce cost* (reduce operational costs, e.g. electricity); *High-margin* (focus on selling high-grade fuel products); *Know market* (having local knowledge of the market); *P. marg. (mis.)* (mentioning high profit margin without further explanation); *Location: Price* (having a location that makes it possible to sustain high price).Panel (B) shows the average cognitive skills of managers mentioning each narrative. Panel (C) excludes managers mentioning causes that are voiced by less than 5 percent of managers. Error bars indicate 95% confidence intervals.

vor high prices, suggesting that understanding competitors leads to a wariness of price cuts. Confidence in the ability to influence fuel sales is associated with a significantly lower probability of thinking that high prices are important for profits, consistent with believing in the efficacy of lower prices. A belief that the optimal price is to match the prices of competitors charging below the price ceiling, a potential sign of neglecting competitor price responses, is associated with a significantly lower probability of believing that price needs to be high to achieve high profits. Consistent with mediation, Figure 5 shows that controlling for all of these measures simultaneously substantially reduces the coefficient on cognitive skills (compare Regression 1 and 2). Cognitive skills remain a strong predictor even after controlling for these measures, however, consistent with capturing a general underlying ability of inductive reasoning that is related to understanding competitor sophistication. Overall, the narrative results are in line with the conceptual framework, in that lower cognitive skills make it less likely that a manager views high prices as beneficial for profits, through a mechanism of understanding competitors and beliefs about the effectiveness of price cuts.

Figure 5: High price narrative, cognitive skills, and mental models of competitors

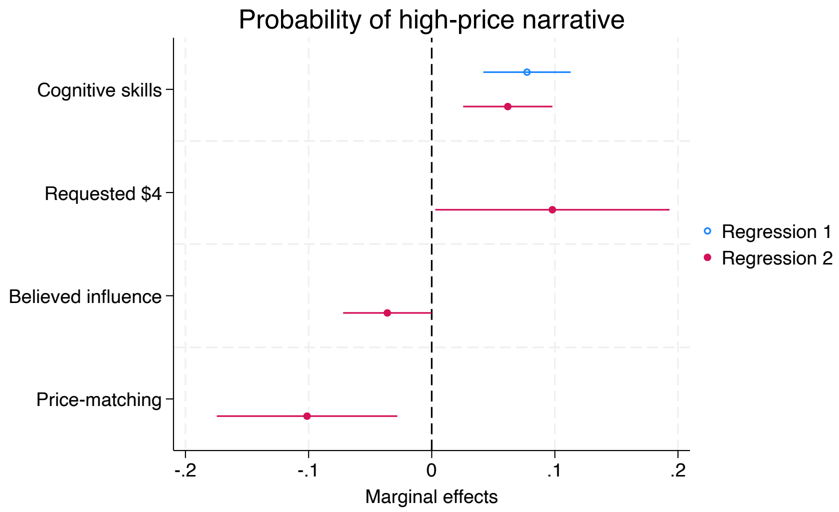

Notes: The figure reports marginal effects from Probit regressions, with 95% confidence intervals. The dependent variable is an indicator for whether a manager mentioned the high price narrative. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, age, gender, experience), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators, district fixed effects). Regression 2 includes these traits but adds three measures of mental models of competitors: requested $4 in the money request game, believed influence over fuel sales, and believed optimality of price-matching strategy. Appendix Table B.3 provides the underlying regression estimates.

Robustness: As detailed in Appendix B.3, we find similar results using alternative methods for reconciling RA disagreements in the classification (e.g., a third RA). Results are also very similar using the sample of managers who mention only one cause, ruling out that the greater likelihood for high skill managers to mention high price is a mechanical outcome of tending to mention multiple causes. A different concern could be that our results reflect correlated measurement error arising from low effort on survey responses, e.g., choosing "don't know" as the easiest responses in the narrative question, and putting in low effort on the Raven test. However, the link between high cognitive skills and mentioning high price remains robust when excluding "don't know" responses and focusing on managers who go to the trouble to cite specific causes. While Figure 5 uses first-wave measures of all independent variables to explain high price narratives, we find similar results using second-wave measures of all available independent variables (Figure B.10). We find a similar relationship between cognitive skills and mentioning high price within different market conditions and location types, indicating that this is not driven by managers with different cognitive skills facing different types of local conditions (Figure B.5)

4 Cognitive skills and pricing behavior

In this section, we turn to investigating managers' preferences and decisions about fuel pricing at their own stations.

4.1 Self-reported pricing preferences, proposal behaviors, and cognitive skills

We begin by analyzing self-reported pricing preferences and self-reported price-change proposal behavior, elicited in our second survey wave. Managers were asked: (1) whether they prefer prices lower than the defaults set by upper management, and (2) if they have to make proposals to change fuel prices in their district, how often they request cuts. As shown in Appendix C.1, managers with lower cognitive skills are significantly more likely to favor lower prices (Figure C.1) and to propose price cuts to upper management (Figure C.2). Among those favoring lower prices, 80% believe this "benefits the company," consistent with the narratives evidence that lower-skill managers view price cuts as helping the company's bottom line.[23] Indeed, Figures C.1 and C.2 show that managers citing high price in their narratives about causes of high fuel profits are significantly less likely to report desiring or proposing price cuts. Self-reported desire for price cuts and tendency to propose price cuts remain significantly related to cognitive skills after controlling for other manager traits, as well as station and market characteristics. Mental models also help explain variation: Requesting $4 in the money request game predicts significantly lower probability of desiring price cuts, and lower frequency of proposing price cuts, while believing in sales influence or optimality of price-matching are significant in the opposite direction (Figures C.3 and C.4).

[23]The question asked managers to indicate whether they "would like to charge lower prices than the default set by upper-level management," and if yes, whether the price cuts benefit themselves, or the company, or both.

4.2 Relationship of actual pricing to manager cognitive skills

Turning to actual pricing, it is clear that pricing decisions can reflect a manager's need to react to many different factors that arise in a changing and noisy environment, e.g., shocks to demand, costs, competitor actions, and senior manager decisions. We investigate whether lower cognitive skills nevertheless affect prices, leading to a systematic tendency to charge lower prices on average. We use our monthly panel data on the pricing behavior of gas station managers. Due to varying price ceilings and the diverse pricing of different fuel products, the company computes a metric to measure a station's overall pricing for a month. This metric compares the monthly average price of each fuel product to its respective ceiling, weighting each product by its sales volume at the station in that month. A ratio of 1 indicates pricing equal to the ceiling for all products. This price ratio incorporates all forms of discounts, including reductions in listed prices and all types of coupons and promotions.

Figure 6: Pricing behavior as a function of cognitive skills and mental models

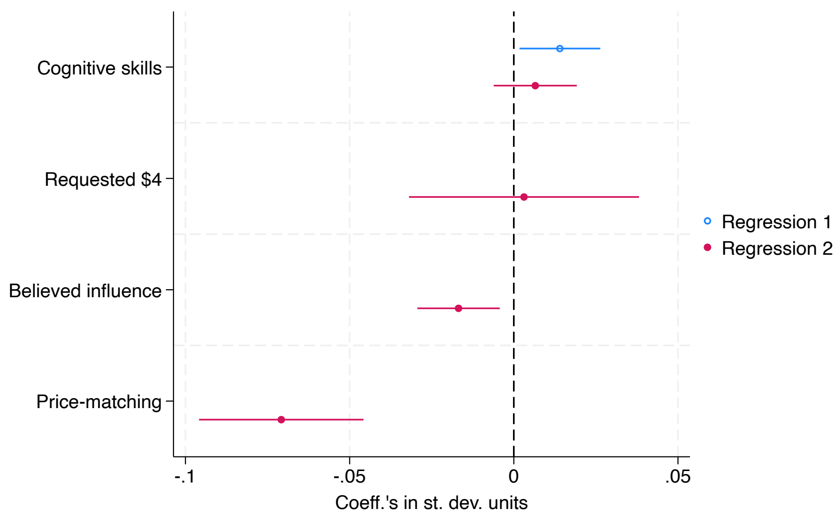

Notes: This figure reports coefficients from OLS regression on monthly price ratios, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the standardized station monthly price ratio relative to the price ceiling. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, age, gender, experience), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators), and interacted month and district fixed effects. Regression 2 includes the three mental models. Appendix Table C.1 provides the underlying regression estimates.

In Figure 6, Regression 1 shows that managers with lower cognitive skills tend to set significantly lower fuel prices on average relative to the ceiling, controlling for other manager traits as well as station and market characteristics and time fixed effects. This is consistent with the conceptual framework, and makes sense in light of our results on how cognitive skills are related to mental models of competitor behavior and narratives about causes of fuel profits. The magnitude of the relationship between cognitive skills and price ratio is substantial in comparison to the reduction in price associated with adding an additional independent competitor, roughly half as large, or the effect of another major competitor station, roughly one-fifth as large (see coefficients in Table C.1 on independent and major competitor, respectively).[24] These magnitudes roughly double if one takes into account the attenuation arising from measurement error in cognitive skills, i.e., the relationship is equal to adding one whole independent competitor or about one third of a major competitor (see discussion of attenuation in Section 2).

Regression analysis shows that mental models of competitors are also related to actual pricing in the expected directions, although less precisely estimated. The indicator for requesting $4 in the money request game is not statistically significant, but beliefs about influencing sales and optimality of price-matching are both significantly negatively related to actual prices (Columns (2) to (4) in Table C.1). Controlling for all three mental model measures simultaneously reduces the coefficient on cognitive skills substantially, by 30%, consistent with these playing a mediating role (compare Regressions 1 and 2 in Figure 6) and the three mental model measures are highly jointly significant (F-test; $p < 0.001$).

4.3 Event study

In this section we present results from an event study design, comparing pricing at stations before and after a new manager arrives. One purpose is to provide an additional robustness check on causality. To be clear, we already have several reasons to believe the results are causal rather than being driven by omitted variables: Observable characteristics of stations and markets do not differ substantially by manager cognitive skills, we control for observable station and market characteristics in our regressions, and we have evidence of plausible causal mechanisms through which cognitive skills would affect pricing decisions, namely different mental models of competitor behavior and optimal pricing strategies. The event study adds further evidence, however, by controlling for *unobservable* station and market characteristics. A second purpose of the event study is to provide a glimpse into the dynamics of how bounded rationality influences pricing strategies over time, e.g., whether the effect is persistent or whether it decreases with feedback.

We analyze the roughly 4,500 manager-change events in our dataset for which we know the traits of the new manager. Since we want to understand how variation in cognitive skills across newly arrived managers affects the time profile of their pricing decisions, we estimate treatment effects separately for each treated station. We want to correct the before-after comparison for time trends in a difference-in-differences analysis, but it is challenging to find a single control station with good pre-trends for every single treated station. We therefore use the synthetic difference-in-differences (SDID) method proposed by Arkhangelsky

[24]The results in Table C.1 show a positive coefficient on number of other stations from our partner company in the local market, consistent with it being easier to sustain high prices when there are more stations from the partner company.

et al. (2022). The intuition builds on the synthetic control method of Abadie et al. (2015): even though no single control station may perfectly match a treated station in terms of characteristics and pre-trends, we can create a synthetic control that combines multiple control stations with appropriate weights to produce an aggregate that mimics the treated station's characteristics.[25] Unlike the original synthetic control method, however, SDID does not require identical levels between treatment and control in the pre-period, only parallel trends. This approach addresses various concerns about non-random assignment of managers to stations that differ on unobservables. For example, if high-skill managers tend to be assigned to stations where it is easy to charge high prices, such factors should lead to high prices in the pre-period, and these are differenced out. See Appendix C.2.1 for a more detailed discussion.

Figure 7: SDID treatment effects, cognitive skills, and mental models

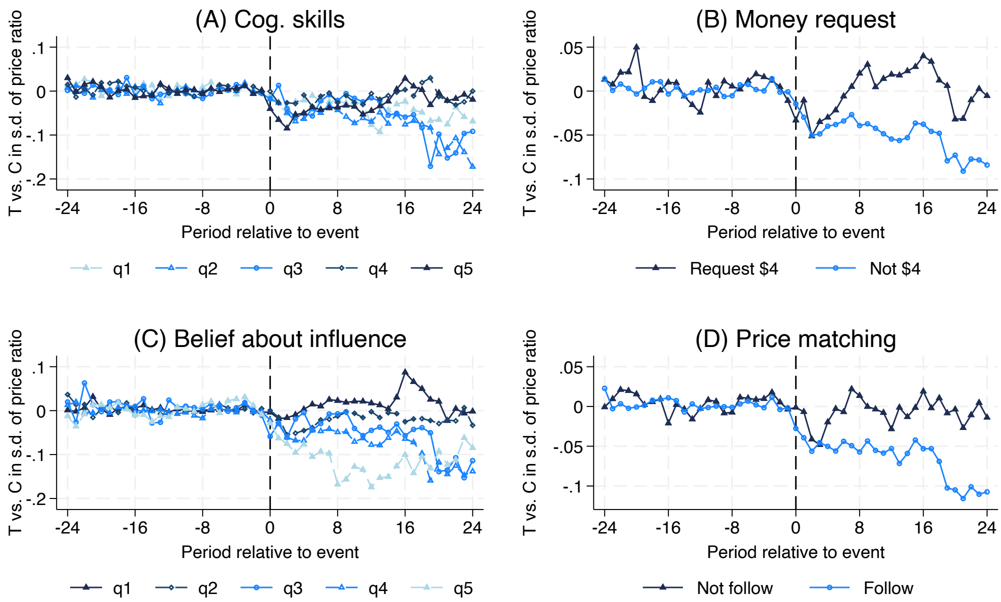

Notes: This figure reports treatment effects on price ratio in the treatment stations versus the synthetic control stations from SDID regressions, by traits of the *new* manager. The treatment stations are stations that went through a manager change, and the synthetic control stations are drawn from stations without a manager change during the same period. The treatment effects are categorized in Panel (A) by quintiles of cognitive skills (quintile 5 is the highest), in Panel (B) by requesting $4 in the money request game or requesting a different amount, in Panel (C) by quintiles of belief in ability to influence fuel sales (quintile 1 is little influence), and in Panel (D) by whether a manager believes that price-matching is optimal.

Panel (A) of Figure 7 reports the average SDID treatment effects on price ratios for each month, depending on the new manager's cognitive skills. In a given month the treatment effect is the difference between the treated station's price ratio and its synthetic control's price ratio in that month, minus the weighted average difference between these two during

[25]For example, if our treated station has pricing patterns that fall between those of a highway station and an urban station, the synthetic control might assign 60% weight to highway stations and 40% weight to urban stations in the control group to best replicate the treated station's behavior.

the pre-treatment period. The figure plots the average of the treatment effects over time, where positive values indicate that the treated station has higher prices than its control in that month, relative to the average pre-period difference. In the pre-period, the estimated treatment effects are close to zero across all cognitive skill quintiles. This confirms that our synthetic controls successfully replicate the pricing behavior of treated stations before the manager change, validating the parallel trends assumption. However, once the new managers arrive (period 0 and beyond), we see changes in pricing, with the evolution depending on cognitive skills. Interestingly, the graphs suggest that, regardless of cognitive skills, new managers tend to reduce prices initially, potentially indicating an experimentation phase. Subsequently, however, managers in the top quintiles of cognitive skills begin raising prices, while those in the lower quintiles reduce prices over time. The divergence becomes pronounced starting at around 12 months and shows no signs of diminishing.

A possible explanation for these dynamics is that the differences in mental models across managers with high and low cognitive skills lead to different explanations for what they observe in their environment. For example, a failure to anticipate or recognize causal effects of own price cuts on competitor prices might lead low-skill managers to react to initial non-success of price cuts by concluding they should reduce prices even more.

Panels (B) through (D) of Figure 7 show a similar analysis, but according to the mental models about competitors of the new manager. We again see a tendency for prices to drop initially with the new manager, but then different evolution of pricing according to mental models. Bringing in a manager who requested $4 in the money request game leads to higher prices over time, compared to managers who do not. Likewise, bringing in a manager who believes he or she can influence fuel sales, or who believes price-matching is optimal, leads to falling prices over time. These effects also take some time to emerge, but the divergence is already apparent after a few months and shows no signs of diminishing even at the end of the two-year post-period.

We reach similar conclusions based on a regression analysis, regressing the SDID treatment effects on traits and mental models of the new manager (see Appendix Figure C.5). The results show a consistent pattern: having a new manager with low cognitive skills is associated with lower prices, with the difference becoming (marginally) statistically significant if we consider 12 months after the change. Similar regressions show that the mental model measures are generally significant or marginally significant, controlling for cognitive skills, even including the periods before 12 months, but point estimates get larger considering time frames 6 months, or 12 months, after the change (Columns (2) to (4) of Appendix Table C.2 to Table C.4). If we include all three mental model measures in the regression simultaneously, along with cognitive skills, they are highly jointly significant (entire post-period sample, F-test; $p < 0.001$), and the coefficient on cognitive skills is reduced, consistent with a mediating role. In summary, the SDID analysis adds further evidence that cognitive skills cause persistent differences in pricing, and this is due in part to how cognitive skills lead to different mental models of competitors.

Robustness: In Appendix C.2.2, we provide various robustness checks. We show that our SDID regression results are robust, and if anything even stronger, if we correct for the fact that the dependent variable is an estimated variable using empirical Bayes shrinkage. We also address another concern, that there could be a special type of time trend for stations receiving a new manager, based on anticipating the arrival of a new manager, which cannot be captured by a synthetic control using only stations that never have a change. As we explain in more detail in the appendix, our SDID analysis is already somewhat robust to this concern, because we use a long pre-period. We also present confirmatory results from an alternative event study design, which considers only stations that receive a new manager, ensuring both groups face similar underlying time trends, and does not need SDID, because we estimate treatment effects at the group level. The difference-in-differences analysis compares pricing in the post-period for the group of above-median cognitive skill new managers (treatment) to the group of below-median cognitive skill new managers (control), and finds a positive difference in price levels over time.

5 Cognitive skills and price wars

To explore whether the impact of bounded rationality on pricing at individual stations spills over into the local market, potentially causing price wars, we use data from one region, which includes around 900 gas stations with daily price information for all fuel products as well as competitor prices in the local market.

To test whether the lower prices charged by low skill managers could trigger price cuts by competitors, we examine how competitor prices change after price-cut events by our partner stations. A price cut event is defined as when our partner station maintains the same price for 7 days and then reduces it by at least 10 cents, thereby moving further below the price ceiling. To estimate a causal effect, we employ a difference-in-differences approach, comparing the price responses of treated competitors who face a price cut event to control competitor stations in the same district whose competing partner stations maintained stable prices during the same period without implementing a price cut.

As shown in Figure 8, competitors are indeed responsive to our partner stations' price cuts. Prior to the price cut, competitor prices show no differential trends. However, immediately following our partner stations' price cuts, treated competitor stations reduce their prices by approximately 5 cents on average, and this price reduction persists over the following week.

Another way that managers with lower cognitive skills could trigger price wars is if they respond more aggressively to competitor price cuts. As we document in Appendix D.1, this is indeed the case, with low-skill managers responding more strongly to a given competitor price cut than high-skill managers. This pattern of price responses suggests that managers with lower cognitive skills may end up in sustained periods of competitive price cutting, namely price wars.

Figure 8: Competitor's price response to partner station price cut

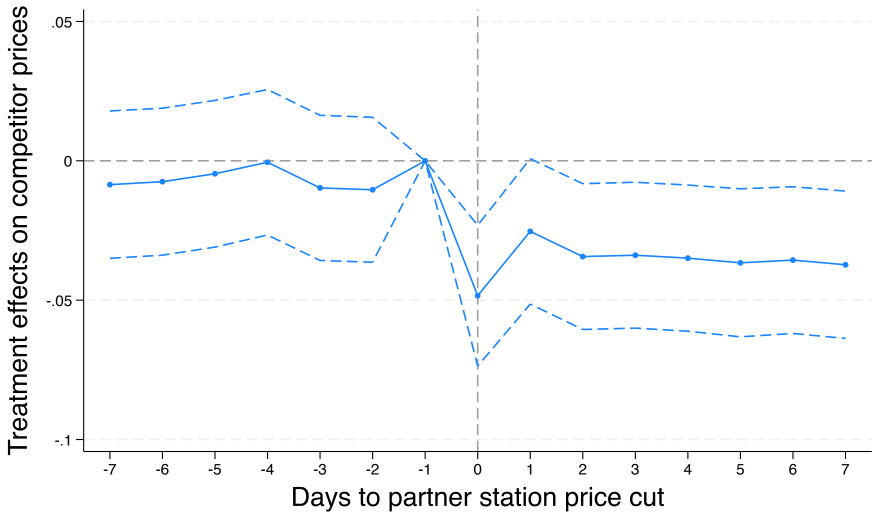

Notes: This figure reports difference-in-differences treatment effects of competitor price responses to partner station price cuts with 95% confidence intervals. A price cut event is defined as when a partner station maintains the same price for 7 days and then reduces it by at least 10 cents. The control group consists of competing stations in the same district where the partner station maintained the same price for 7 days during the same period but did not implement a price cut. The solid line shows point estimates of treatment effects relative to the price cut event (day 0).

To analyze whether managers with lower cognitive skills do end up more frequently in price wars, we need to define a price war. There is no accepted quantitative definition in the literature, which is understandable because the magnitude of price cuts that constitutes a war is clearly context-dependent. A working qualitative definition proposed by Busse (2002) emphasizes that price wars involve cuts "significantly below the usually prevailing prices." We define a price war as *mutual price cuts of at least 50 cents from the price ceiling for a period of 14 days or more.* Here "mutual price cuts" mean a partner station and at least one competitor station in the local market were involved in the war. For an example of a price war in our data, see Figure D.2. We observe 211 price wars in 1324 non-monopoly fuel markets between 2018 and 2021. On average, a price war lasted 43 days (median 29 days) and our partner stations lowered their prices by 68 cents relative to the price ceiling during the price war period.

To analyze how the tendency to be involved in price wars relates to manager cognitive skills, we aggregate to the manager level, with the dependent variable being the number of price wars in which the manager is involved. Figure 9 plots the relationship between managers' cognitive skills and their tendency to end up in price wars. While managers in the lowest cognitive skill quintile had around 0.3 price wars in the three-year period, the number of wars steadily decreased to less than 0.15 for managers in the highest cognitive skill quintile. Roughly speaking, managers with the lowest cognitive skills are engaged in price wars about twice as often as managers with the highest cognitive skills. The wide confidence intervals for individual quintiles reflect the fact that this is a part of our analysis where the sample size is relatively small; since we use a single region and around 70% of managers in this region answered the survey, we have approximately 650 observations.

Figure 9: Cognitive skills and price wars

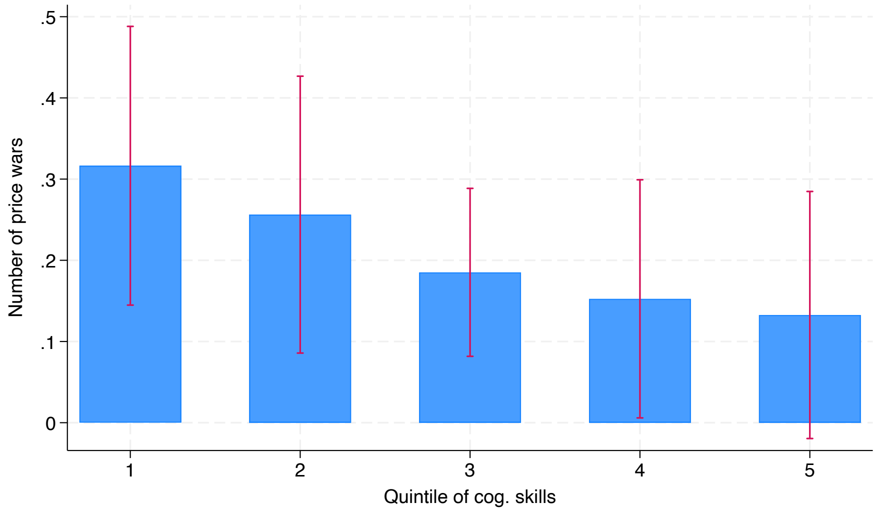

Notes: This figure reports the relationship between cognitive skills and the number of price wars between July 2018 and January 2021. The horizontal axis shows quintiles of cognitive skills, with 1 being the lowest and 5 being the highest. The vertical axis represents the average number of price wars for managers in each cognitive skill quintile. The 95% confidence intervals for the mean number of price wars are shown as error bars.

Regression analysis (Regression 1 in Appendix Figure D.3) shows that the number of price wars is significantly negatively related to manager cognitive skills, controlling for other manager traits as well as station and market characteristics. The regression also controls for days of operation to correct for the mechanical tendency for stations to have more price wars if they operate longer.

We also regress the number of price wars on our various measures of mental models of competitors. Point estimates are in the expected directions (Columns (2) to (4) of Table D.1), but estimates are imprecise. When we control for all three measures simultaneously, the coefficient on cognitive skills is reduced, and belief in ability to influence fuel sales becomes a marginally significant positive predictor of number of price wars (compare Regressions 1 and 2 in Figure D.3).

In conclusion, while some price wars may be optimal actions that are part of a strategy to sustain high prices, the fact that these are substantially more common among managers with low cognitive skills suggests that at least some may instead be strategic mistakes.

Robustness: Results on competitor price responses are robust to other definitions of a price cut event (Figures D.4 and D.5), and to generating the control stations for each price cut event using SDID approach (Figure D.6). Regarding price wars and cognitive skills, we find similar results using an alternative, less-conservative definition involving a price cut of 25 cents (Figure D.8). We also check whether interdependence of local markets could affect our results (price wars could be correlated across nearby markets, because they share a competitor). We find similar results in a robustness check focusing only on "isolated markets," defined as stations that do not share any competitors with other partner stations (see Figure D.7).

6 Additional evidence on mental models of competitor pricing

Our findings so far are consistent with the idea that managers with lower cognitive skills are less likely to take into account competitor responses to price cuts. To gather even more direct evidence, however, we designed a prediction task in our third survey wave, which we implemented for a randomly selected sub-sample of roughly 6,000 managers.[26]

As shown in Table 3, managers were presented with real historical data from a partner station. The station initially priced at the ceiling (5.01 per liter in local currency), while being undercut by two independent stations selling inferior brands; the partner station later reduced its price to 4.81, which was followed by a further price reduction by one competitor (Competitor 1) after a lag of a few days, while the other competitor (Competitor 2) did not respond. We asked managers to predict sales and competitor prices if the station had counterfactually charged 4.51 instead of 5.01 during the initial period. Their predictions provide an indication of whether they anticipate competitor price responses to this earlier and relatively deep price cut.

We observe systematically different predictions of competitor prices depending on managers' cognitive skills. As shown in Figure 10, only 29% of managers in the bottom quintile of cognitive skills predict that Competitor 1 would respond to the hypothetical price cut by lowering their prices. This share increases with cognitive skills, reaching 46% for managers in the top quintile. We also observe a similar difference for predictions of Competitor 2's price (Figure E.2). Regression analysis shows that the probability of predicting a competitor price reduction increases significantly with cognitive skills (Regression 1 in Figure E.3), controlling for other manager traits as well as station and market characteristics.

A closer examination of the complete distribution of predicted prices provides additional insights (Appendix Figure E.1). While the majority of managers with low cognitive skills predict that the price of Competitor 1 would stay the same when our partner station lowers its price, those who do predict changes are roughly as likely to predict price increases as

[26]The other managers were presented with a different set of measures, to be used for another study.

Table 3: Historical Data and Hypothetical Scenario Presented to Managers

	(A) Actual Historical Data				
Date	Station price	Sales	Competitor 1 price	Competitor 2 price	Profit
May 11	5.01	2,118	4.28	3.98	2,118.5
May 12	5.01	1,981	4.28	3.98	1,980.8
May 13	5.01	855	4.28	3.98	854.6
May 14	5.01	1,530	4.28	3.98	1,530.3
May 15	5.01	3,334	4.28	3.98	3,334.0
May 16	4.81	1,600	4.28	3.98	1,280.1
May 17	4.81	2,841	4.28	3.98	2,272.7
May 18	4.81	1,528	4.29	3.98	1,222.7
May 19	4.81	1,764	4.29	3.98	1,411.1
May 20	4.81	4,279	4.29	3.98	3,423.5
May 21	4.81	679	4.08	3.98	543.5
May 22	4.81	2,014	4.08	3.98	1,611.2
May 23	5.01	2,179	4.08	3.98	2,179.3
May 24	4.81	2,203	4.08	3.98	1,762.2
	(B) Hypothetical Scenario				
Date	Station price	Sales	Competitor 1 price	Competitor 2 price	
May 11	5.01	2,118	4.28	3.98	
May 12	4.51	_____	_____	_____	
May 13	4.51	_____	_____	_____	
May 14	4.51	_____	_____	_____	
May 15	4.51	_____	_____	_____	

Notes: Top panel reports actual historical data of a station of our partner company. Bottom panel presents the counterfactual scenario where the station reduced prices on four days. Blanks (_____) indicate fields managers were asked to predict. Managers were presented with the following description: "The following is the real price and sales volume data at a gas station of your company within 14 days. During these 14 days, the price ceiling is kept constant at 5.01 per liter. It can be seen from the table that this gas station has two competitors, both of which are independent gas stations. The second column shows the price of the own-company gas station. The third column shows the daily sales volume at this gas station. The fourth and fifth columns show the prices of the first and second competitors respectively. The sixth column shows the profit at this own-company gas station."

price decreases. This suggests noise, or uncertainty about the causes of competitor prices. In contrast, managers in the top cognitive quintile almost never predict price increases, have the lowest rate of predicting a constant price, and most predict a decrease.[27]

Consistent with the idea that predicting a competitor price decrease indicates an underlying ability to understand competitors, managers who requested $4 in the money request game were substantially more likely to predict competitor price reductions (43% vs. 34%). In a regression analysis, adding the indicator for requesting $4 to the regression on cognitive skills reduces the coefficient of cognitive skills (Regression 2 in Figure E.3), and requesting $4 is associated with a significantly higher probability of predicting a competitor price reduction.

[27]Our measure may provide a lower bound on whether managers think competitors respond to price cuts with price reductions, because the data showed a five-day lag in price response, whereas managers were asked about predicting competitor prices over a four-day window. If managers think the timing shown in the (real) data reflects a strict rule of the competitor to wait five days, they might predict a constant price.

Figure 10: Cognitive Skills and Predicted Competitor Price Responses

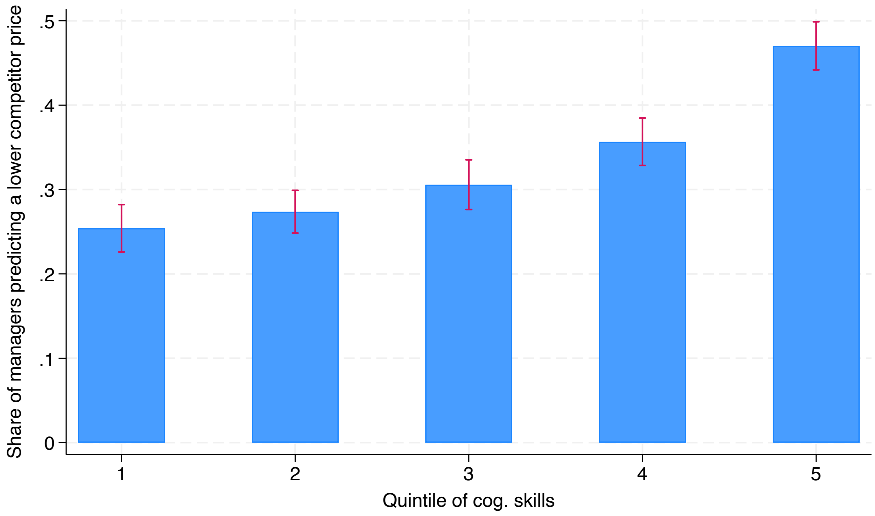

Notes: This figure reports the share of managers predicting a lower competitor price in response to a hypothetical price cut of a partner station by quintiles of cognitive skills (quintile 5 is the best).

The results also support our interpretation of beliefs about ability to influence fuel sales, and about the optimality of price matching, as outcomes of failure to predict negative competitor price responses. We regress each of these on the indicator for predicting a negative price response and find that managers who predict a negative price response are significantly less confident in their ability to influence fuel sales and less likely to think that price-matching is optimal, controlling for cognitive skills and other manager and market characteristics (Figure E.4).

In summary, managers with lower cognitive skills appear less likely to perceive a causal link between their own pricing decisions and competitor responses, even when presented with identical data showing such interactions. This can help explain why these managers favor low prices and more frequently engage in aggressive price cutting—they systematically underestimate how competitors will respond to their actions.

7 Consequences for profits, welfare, and market power

7.1 Profit consequences of price cuts

Several pieces of evidence indicate that the more aggressive pricing of managers with low cognitive skills may, in fact, be a mistake that contributes to low profits. A first observation is that, since cognitive skills are related to the ability to predict competitor behavior and are a measure of decision quality, there is already reason to think that pricing strategies associated with low cognitive skills may be less beneficial.

A second observation is that total profits are positively correlated with higher fuel price, in a regression of profit on price, controlling for manager traits and station and market characteristics (Column (2) of Table F.1). This is driven by a positive relationship between fuel prices and fuel profits (Column (5)). There is a negative relationship of fuel prices to convenience store profits, which makes sense given that low fuel prices can attract more customers to the store (Column (8)). The relationship to convenience stor profits relationship is relatively weak, however, and fuel profits make up the bulk of total profits, explaining why the overall relationship to total profits is negative.

An issue with regressing profit on price is the possibility of various sources of endogeneity. For example, there could be reverse causality if low profits cause managers to either increase or decrease price. As a way to taddress this we use the measures of mental models of competitors – requesting $4 in the money request game, belief about ability to influence fuel sales, and belief about the optimality of price matching – to instrument for the price ratio. We have shown in Figure 6 that the mental models are highly jointly significant in explaining the monthly price ratio, providing a strong first stage. In addition, all three measures are relatively narrowly focused on understanding competitor strategic behavior and beliefs about optimal pricing strategy, which increases the plausibility that these influence profit only through the channel of pricing (the exclusion restriction).

Figure 11: Profits as a function of instrumented price

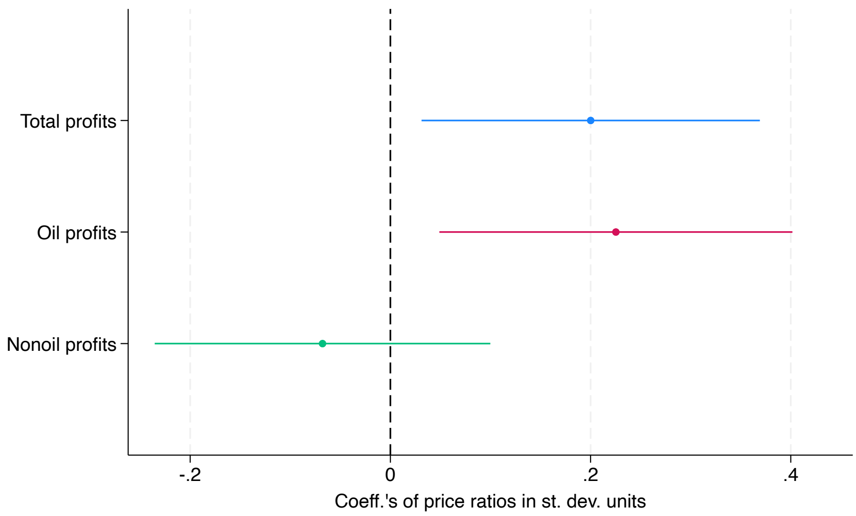

Notes: This figure reports coefficients from 2SLS regression on standardized monthly profits, instrumenting standardized price with mental models, with 95% confidence intervals based on robust standard errors clustered at the station level. Controls include manager traits (cognitive skills, noncognitive skills, experience, gender, and age), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators, and market share), and interacted month and district fixed effects. The mental models include winning in the money request game, the price-matching strategy, and the confidence in influencing fuel profits. Appendix Table F.1 provides the underlying regression estimates.

Figure 11 provides a third piece of evidence, showing the results of two-stage least squares regressions, explaining station performance with instrumented price ratio. We see that total profits are significantly positively related to price ratio. This relationship is mainly driven by higher fuel profits. The point estimate is negative for convenience store profits, consistent with low fuel prices increasing nonfuel profits, but this is not statistically significant.

A potential concern could be that our analysis of profit consequences neglects some other types of benefits of the price cuts by low-skill managers. For example, perhaps there could be long-run profit benefits, e.g., if price cuts discipline competitors and help foster high prices and profits in the future. The fact that low-skill managers end up in price wars more often, however, seems contrary to the idea that low-skill managers are better at sustaining high prices through price wars. Furthermore, our four-year data on station performance show that managers with lower cognitive skills earn lower average profits (see Column (1) in Table F.1), providing no evidence that the low-price strategy of low-skill managers yields long-term advantages. Another idea might be that the partner firm places a value on high sales, independent of current profits. This is contrary to our understanding of company objectives based on discussions with managers, however, and as we discuss next, our survey of senior management provides emphatic evidence of a widespread concern about sub-optimally low prices.

Our survey of district-level managers provides a final source of evidence that price cuts of managers can be sub-optimal. As shown in Table 4, it is a pervasive view of district managers that station managers have a tendency to cut prices, if they are given more autonomy over price. Furthermore, when asked whether the price chosen by station managers would be too high, about right, or too low, roughly 75% say too low. In an open-ended question, we also asked district managers whether station managers should be "allowed to actively participate in price competition and charge lower prices than competitors." Classifying responses, we find that 74% of district managers indicated that such strategies are not always beneficial. Among these, 71% took a stronger position, stating that aggressive price competition is always a bad idea. When examining the reasoning provided by those who viewed price competition as always harmful and offered specific explanations, we found their concerns fell into two main categories: half stated that lower prices directly lead to lower profits, while the other half explicitly mentioned the risk of triggering price wars with competitors. These findings indicate that district managers see causality going from price cuts to low profits, and suggests that they do not see benefits of a low price strategy in terms of long-run profits or because of an intrinsic or long-run value of increasing sales.

7.2 Welfare consequences

To calculate the magnitude of the welfare effects of boundedly rational pricing, we need to make additional assumptions about the nature of demand for fuel and how this depends

Table 4: Views of district level managers about the optimality of price cuts

If station managers have full autonomy over price setting, do you think the price will be:	Frequency	Percentage
Higher than the current price	11	3%
Same as the current price	91	27%
Lower than the current price	236	70%
If station managers have full autonomy over price setting, do you think the price will be:		
A price that is too high	28	8%
The optimal price	53	16%
A price that is too low	257	76%

Notes: Results are from the survey with the district level managers.

on price. To match our conceptual framework, and the evidence we have found on the important role for competitor price responses in the fuel market (Section 5), we want a model in which lower prices can directly affect demand, but also triggers competitor price cuts. We adopt the conjectural variation framework for differentiated products (e.g., Slade, 1984). We provide further details in Appendix F.2.

Within this framework, residual demand for a given firm is assumed to be approximated by the linear equation:

$$q_i = f_i(p_i, p_{-i}, z) = a + b_i p_i + c_i p_{-i} + g(z)$$

where q_i is quantity of fuel (gasoline and diesel) products sold at station i, p_i is own price, p_{-i} is (average) competitor price, and $g(z)$ are demand shifters. To estimate how price changes affect producer and consumer surplus, we need the response of demand to price, $\frac{dq_i}{dp_i} = b_i + c_i\theta_i$ (recall that θ_i is the competitor price response). We use two exogenous cost shifters—the government price ceiling in the *current* pricing cycle, $Ceiling_t$, and the ceiling from the *previous* cycle, $Ceiling_{t-1}$—to estimate own-price response b_i and cross-price effect c_i. Because our partner company's internal accounting rule links marginal cost one-for-one to the current ceiling, $Ceiling_t$ moves our partner stations' prices close to one-for-one, whereas independent competitors adjust only partially and with delay to changes in the world oil price (to which the ceiling is indexed), so $Ceiling_{t-1}$ still influences their prices in period t. We therefore treat $Ceiling_t$ and $Ceiling_{t-1}$ as a pair of excluded instruments for the two endogenous variables p_{it} and p_{-it}. Controlling for station-fuel-product fixed effects, as well as day-of-the-week, month, and year fixed effects for each fuel product, we estimate the own-price effect for an average station $\bar{b}\,\frac{\bar{p}}{\bar{q}}$, which implies an elasticity of -2.07, and we estimate the cross-price effect $\bar{c}\,\frac{\bar{p}}{\bar{q}}$, which implies an elasticity of 1.15. Finally, through event studies of actual price cuts, we estimate the average competitor price response $\bar{\theta}$, finding that competitors match 38% of a price cut. We can then calculate the needed derivative of quantity with respect to own price using the equation $\frac{d\bar{q}}{d\bar{p}} = \bar{b} + \bar{c}\bar{\theta}$. The calculation implies an average elasticity of -1.63. Detailed estimation methods and results are presented in Appendix F.1.

With these parameter estimates in hand, we can now calculate the welfare implications of different pricing strategies for an average station. The impact on producer surplus is the difference in $p_j \cdot q_j - MC \cdot q_j$ for the high- and low-skill managers, with $j \in \{h, l\}$. Since we do not directly observe marginal costs, we perform calculations across a range of plausible values.[28] A 1 s.d. decrease in cognitive skills reduces producer surplus (PS) provided by the station by 2.9% per year given the midpoint value for marginal cost. The loss is in the range of 1.7% to 6.8% for the full range of possible marginal costs. This structurally estimated range brackets our finding using observed profits with the instrumental approach—that the profit loss for a 1 s.d. decrease in cognitive skills is 3%.

While the lower price charged by the low-skill managers reduces producer surplus, it leads to higher consumer surplus. The change in consumer surplus (CS) is measured by the area to the left of the demand curve between the higher and lower prices implied by the difference in cognitive skills. We calculate an increase in CS of 1.3% from having a manager with 1 s.d. lower cognitive skills. Notably, CS is roughly 5 times larger than PS in absolute terms, so a 1.3% increase in CS more than offsets the decrease in PS and implies an increase in total surplus. Specifically, the lower price implies a substantial reduction in deadweight loss, ranging from 4.5% to 14.3% for the range of possible marginal costs.

Measured market power Turning to standard measures of market power based on the markup of price over marginal cost ($\frac{p-MC}{p}$), we can assess the impact of cognitive skills under plausible assumptions about marginal cost. We calculate that a 1 s.d. reduction in cognitive skills reduces the markup by 3 percent (0.29 percentage points) for the midpoint value of marginal cost. The difference ranges from 2.0 to 7.1 percent (0.27 to 0.30 percentage points) for the full range of marginal costs. This means that the same gas station facing the same market conditions can have substantially more or less measured market power depending on the cognitive skills of the manager. Thus, having a lower-skilled manager reduces producer surplus and reduces the measured markup, but at the same time improves consumer surplus and market efficiency.

Eliminating bounded rationality As a counterfactual, we consider the welfare implications of eliminating managerial bounded rationality entirely. We model this as raising all managers' cognitive skills by 2 standard deviations, effectively moving them from the current mean to the 97th percentile level—equivalent to having only the most cognitively skilled managers operate stations. This scenario could represent the introduction of algorithmic pricing systems that match the decision quality of the highest-performing human managers. Under this counterfactual, producer surplus would increase by 3.4% to 13.4% per station for the range of plausible marginal costs, while consumer surplus would decrease by 2.6%.

[28]We estimate the range of plausible marginal costs based on observed wholesale refined oil prices for independent stations (as a lower bound) and allow for up to 60 cents higher costs to account for the premium quality of our partner company's fuel (as an upper bound). See Appendix F.2 for details.

Deadweight loss would increase by 9.0% to 28.6%, and measured market power (markup) would increase by 4.0% to 14.2%. Given that our partner company and the other major firm together control approximately two-thirds of the market, and assuming similar cognitive skill distributions among their managers, the market-wide implications would be substantial. If both major firms adopted such systems, the aggregate effect would be a significant reduction in market efficiency, with producers capturing additional surplus at consumers' expense.

8 Why are boundedly-rational managers allowed to set prices?

A natural question given our findings is, why does upper-level management still give station managers the autonomy to set prices: Does this reflect bounded rationality on the part of upper-level management, or is it an outcome of a trade-off faced by upper-level management?

Regarding upper-level management's awareness, our district manager survey has shown that many senior managers are concerned about station managers setting sub-optimally low prices (Table 4). There is also some indication that this concern translates into restrictions on station managers' autonomy; we find that districts where senior managers believe station managers would set sub-optimally low prices are significantly more likely to require proposals for fuel price changes compared to districts without this concern (55 versus 41%; Wilcoxon tests, $p < 0.05$). Also, when station managers must submit a proposal to change fuel prices, 66% of these proposals are rejected according to district managers. Even when station managers have the autonomy to change prices without prior approval, they sometimes face pre-specified price ranges within which they must operate. Thus, upper-level management clearly recognizes potential issues with managers' low-price strategies.

Given that they are aware of the issue, the next question is why district-level managers do not completely centralize pricing. One reason is strategic considerations. In our survey of district managers, we asked whether it is a good idea to allow station managers to "actively engage in price competition and charge lower than competitors." Among district managers who supported this approach, one explicitly stated that giving station managers no autonomy and committing to a price would allow competitors to undercut and capture the entire market. By contrast, allowing station managers the possibility to respond could help regulate competitors' pricing behavior. This perspective suggests that some autonomy in pricing serves as a credible threat that deters competitors from aggressive price cutting.

A second reason for granting autonomy is that senior managers view local managers as possessing valuable local knowledge. When asked "where greater autonomy would be beneficial," 78% of district managers cite locations where managers' local knowledge is important, followed by stations with competent managers. Plausible types of local knowledge include understanding the strategies of local competitors, and the tastes of local consumers. Without such knowledge, it is more difficult to know the optimal price at a given station. Indeed, as shown in Appendix F.1, if we estimate proxies for station-specific measures of elasticity, these range from near 0 to -5. Simple uniform pricing rules, such as always pricing at the price ceiling, are therefore unlikely to be optimal (see also DellaVigna and Gentzkow, 2019, for evidence that uniform pricing is not optimal for retail in general). At the same time, as we discuss in the appendix, these elasticity estimates do not allow pinpointing the optimal average price for each station. Identifying competitors' strategies in a station's local market, e.g., how responsive competitors are to price cuts, would also be needed for optimal pricing, but our data do not include enough price cut events to provide credible station-level estimates. These empirical challenges also affect senior managers, reinforcing why senior management values station managers' local knowledge about competitor strategies and customer characteristics.

9 Conclusion

This paper adds nuance to the standard assumption of fully rational firms in price competition by documenting systematic heterogeneity in behavior based on cognitive skills. Higher-skill managers better anticipate competitor reactions, maintain high prices at focal points conducive to coordination, and achieve higher profits—patterns broadly consistent with standard theory. In contrast, lower-skill managers are less aware of competitor sophistication, less likely to view high prices as profit-maximizing, more likely to charge lower prices, more likely to engage in price wars, and earn lower profits. That bounded rationality has substantial effects on pricing and outcomes is particularly striking given the high stakes, learning opportunities, and potential institutional safeguards. Even less obvious is the directional nature of the effect—leading managers to charge lower prices—which may enhance market efficiency.

The mechanisms we document are likely to extend beyond the specific firm, market, and country we study. Our results show that bounded rationality works through weakening perceived competitor price responses, and the comparative static that this will lead to lower prices seems quite general, applicable to any price competition where punishments are important, not just retail fuel markets. It is also noteworthy that our measure of bounded rationality, the Raven test, is highly abstract and does not rely on context specific knowledge about fuel markets. While a distinctive feature of our setting is the presence of a price ceiling, our results on mechanisms suggest that this is not necessary for bounded rationality to depress prices. If bounded rationality simply introduced noise, ceilings might mechanically push prices down by constraining upward errors. Instead, we find evidence of systematically biased beliefs about competitor behavior, which would influence pricing even in markets without ceilings.[29] Our results also seem relevant for a range of organizational structures. They show how bounded rationality can distort pricing even in large firms that delegate to local decision makers. While some firms may avoid delegation to limit pricing errors, doing so can lead to inefficiencies by ignoring valuable local knowledge, and this is itself an important impact of bounded rationality. The role of bounded rationality is likely even more pronounced in small, owner-operated firms, which represent a substantial share of global business activity.

The findings have a number of important implications. They highlight the value of incorporating bounded rationality into theories of strategic price competition, and they provide some insights into how this might be done. They also show how markets may be more efficient than predicted by standard theory, because bounded rationality fosters price competition and counteracts the market failure of market power. Our results also offer guidance for competition policy. They suggest that, while some price wars may signal optimal strategies to sustain high prices, others may simply reflect competitive behavior of boundedly rational managers. Standard measures of market power that rely on price markups may be biased by variation in cognitive skills among local managers. Finally, our calculations imply that introducing algorithmic pricing, if this mimics the most skilled managers in our data, may tend to substantially increase market prices and reduce market efficiency by eliminating the human tendency toward excessive competition that we document.

While our primary focus is on cognitive skills, along the way our analysis provides some results on how other manager traits relate to mental models, pricing, and profits. Success in the money request game is unrelated to manager traits besides cognitive skills. Noncognitive skills and experience predict significantly higher confidence in ability to influence fuel sales, and in the optimality of price matching, so these actually have the same influence as low cognitive skills (Table B.2). On the other hand, noncognitive skills and experience predict a greater likelihood of mentioning high prices in profit narratives (Table B.3). Since these traits do not lead to more accurate mental models of competitor behavior, mentioning high prices seems to reflect some other way of thinking, for example, recognizing that high prices can be beneficial without being certain of how to achieve this. This interpretation is corroborated by the finding that these traits are unrelated to, or, in the case of experience, even negatively correlated with actual prices (Table C.1). These findings suggest some influences of these traits on mental models and pricing, but not the same robust and consistent effect as cognitive skills. This does not mean, of course, that such traits are not important for other aspects of being a good manager.

The findings in this paper raise questions for future research. For example, it would be

[29]Although price ceilings may facilitate coordination on high prices, prior research shows collusion arises in markets without ceilings as well, leaving room for bounded rationality to undermine market power and high prices in such settings. Recent evidence suggests that such coordination often emerges gradually (Byrne and De Roos, 2019); our findings offer one possible explanation for why: Bounded rationality among market participants.

interesting to understand in more detail what may keep boundedly rational managers from learning to charge high prices. Our event study shows that, when facing the same market conditions and experiences, high-skill and low-skill managers seem to learn different things over time. We have suggested that this may reflect bounded rationality leading to a mis-specified mental model in which competitor price cuts are not understood as causal reactions to own price cuts. Our third survey wave provides supporting evidence: low-skill managers are less likely to predict competitor retaliatory price cuts despite observing the exact same data as the high skill managers. Future research could explore in greater depth the structure of the alternative mental model of competitor behavior held by low-skill managers, to understand more precisely what managers think they are learning and why they discount advice from senior management.[30] Additionally, bounded rationality might interfere with backward-looking learning, where managers examine historical data to identify successful pricing patterns. Cognitive constraints could impair this process through biased memory, difficulty analyzing data systematically, or misattributing profit changes to non-pricing factors.[31] Future research could investigate whether cognitive constraints limit managers' ability to learn optimal pricing from their own performance data.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Comparative politics and the synthetic control method," *American Journal of Political Science*, 2015, *59* (2), 495–510.

Abeler, Johannes, David B Huffman, and Collin Raymond, "Incentive complexity, bounded rationality and effort provision," *conditionally accepted at American Economic Review.*

Adhvaryu, Achyuta, Anant Nyshadham, and Jorge Tamayo, "Managerial quality and productivity dynamics," *The Review of Economic Studies*, 2023, *90* (4), 1569–1607.

Alaoui, Larbi and Antonio Penta, "Endogenous depth of reasoning," *The Review of Economic Studies*, 2016, *83* (4), 1297–1333.

Andre, Peter, Ingar Haaland, Christopher Roth, and Johannes Wohlfart, "Narratives about the Macroeconomy," *forthcoming in Review of Economic Studies*, 2023.

[30]Anecdotally, one manager described being forced to abandon a price-cutting strategy that lost money, but insisted the problem was senior management's failure to wait long enough for the benefits to materialize. The manager thus learned that not enough time had elapsed, rather than something else about the problems with price cuts. Such models are hard to falsify, and low-skill managers may be prone to stick to these in the face of advice.

[31]Indeed, Arunachaleswaran et al. (2024) show that near monopoly prices can emerge through backward-looking "no-regret" learning algorithms, suggesting bounded rationality may make managers less effective at such learning than algorithms.

— , **Philipp Schirmer, and Johannes Wohlfart**, "Mental models of the stock market," 2023.

Angrist, Joshua D, Peter D Hull, Parag A Pathak, and Christopher R Walters, "Leveraging lotteries for school value-added: Testing and estimation," *The Quarterly Journal of Economics*, 2017, *132* (2), 871–919.

Aoyagi, Masaki, Guillaume R Fréchette, and Sevgi Yuksel, "Beliefs in repeated games: An experiment," *American Economic Review*, 2024, *114* (12), 3944–3975.

Arad, Ayala and Ariel Rubinstein, "The 11–20 money request game: A level-k reasoning study," *American Economic Review*, 2012, *102* (7), 3561–3573.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic difference-in-differences," *American Economic Review*, 2021, *111* (12), 4088–4118.

Arunachaleswaran, Eshwar Ram, Natalie Collina, Sampath Kannan, Aaron Roth, and Juba Ziani, "Algorithmic collusion without threats," *arXiv preprint arXiv:2409.03956*, 2024.

Asker, John, Chaim Fershtman, and Ariel Pakes, "Artificial intelligence, algorithm design, and pricing," in "AEA Papers and Proceedings," Vol. 112 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2022, pp. 452–456.

— , — , **and** — , "The impact of artificial intelligence design on pricing," *Journal of Economics & Management Strategy*, 2024, *33* (2), 276–304.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market," *forthcoming in Journal of Political Economy*, 2023.

Bandiera, Oriana, Andrea Prat, Stephen Hansen, and Raffaella Sadun, "CEO behavior and firm performance," *Journal of Political Economy*, 2020, *128* (4), 1325–1369.

Baron-Cohen, Simon, Sally Wheelwright, Jacqueline Hill, Yogini Raste, and Ian Plumb, "The "Reading the Mind in the Eyes" test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism," *Journal of child psychology and psychiatry*, 2001, *42* (2), 241–251.

Barron, John M, John R Umbeck, and Glen R Waddell, "Consumer and competitor reactions: Evidence from a field experiment," *International Journal of Industrial Organization*, 2008, *26* (2), 517–531.

Bastianello, Francesca and Paul Fontanier, "Expectations and learning from prices," *Review of Economic Studies*, 2024, p. rdae059.

Berry, Steven, Martin Gaynor, and Fiona Scott Morton, “Do increasing markups matter? Lessons from empirical industrial organization,” *Journal of Economic Perspectives*, 2019, *33* (3), 44–68.

Bilker, Warren B, John A Hansen, Colleen M Brensinger, Jan Richard, Raquel E Gur, and Ruben C Gur, “Development of abbreviated nine-item forms of the Raven’s standard progressive matrices test,” *Assessment*, 2012, *19* (3), 354–369.

Bloom, Nicholas and John Van Reenen, “Measuring and explaining management practices across firms and countries,” *The quarterly journal of Economics*, 2007, *122* (4), 1351–1408.

_ , Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, “Does management matter? Evidence from India,” *The Quarterly Journal of Economics*, 2013, *128* (1), 1–51.

_ , Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen, “What drives differences in management practices?,” *American Economic Review*, 2019, *109* (5), 1648–1683.

Bó, Pedro Dal and Guillaume R Fréchette, “Strategy choice in the infinitely repeated prisoner’s dilemma,” *American Economic Review*, 2019, *109* (11), 3929–3952.

Boissiere, Maurice, John B Knight, and Richard H Sabot, “Earnings, schooling, ability, and cognitive skills,” *The American Economic Review*, 1985, *75* (5), 1016–1030.

Burnham, Terence C, David Cesarini, Magnus Johannesson, Paul Lichtenstein, and Björn Wallace, “Higher cognitive ability is associated with lower entries in a p-beauty contest,” *Journal of Economic Behavior & Organization*, 2009, *72* (1), 171–175.

Busse, Meghan, “Firm financial condition and airline price wars,” *RAND Journal of Economics*, 2002, pp. 298–318.

Byrne, David P and Nicolas De Roos, “Learning to coordinate: A study in retail gasoline,” *American Economic Review*, 2019, *109* (2), 591–619.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, “Artificial intelligence, algorithmic pricing, and collusion,” *American Economic Review*, 2020, *110* (10), 3267–3297.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, “A cognitive hierarchy model of games,” *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Carpenter, Jeffrey, Michael Graham, and Jesse Wolf, “Cognitive ability and strategic sophistication,” *Games and Economic Behavior*, 2013, *80*, 115–130.

Carpenter, Patricia A, Marcel A Just, and Peter Shell, "What one intelligence test measures: a theoretical account of the processing in the Raven Progressive Matrices Test.," *Psychological review*, 1990, *97* (3), 404.

Cattell, Raymond Bernard, *Intelligence: Its structure, growth and action*, Vol. 35, Elsevier, 1987.

Cawley, John, James Heckman, and Edward Vytlacil, "Three observations on wages and measured cognitive ability," *Labour economics*, 2001, *8* (4), 419–442.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and taxation: Theory and evidence," *American economic review*, 2009, *99* (4), 1145–1177.

Corts, Kenneth S, "Conduct parameters and the measurement of market power," *Journal of Econometrics*, 1999, *88* (2), 227–250.

Costa-Gomes, Miguel, Vincent P Crawford, and Bruno Broseta, "Cognition and behavior in normal-form games: An experimental study," *Econometrica*, 2001, *69* (5), 1193–1235.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in US retail chains," *The Quarterly Journal of Economics*, 2019, *134* (4), 2011–2084.

Deneckere, Raymond, "Duopoly supergames with product differentiation," *Economics Letters*, 1983, *11* (1-2), 37–42.

Dimmock, Stephen G, Roy Kouwenberg, Olivia S Mitchell, and Kim Peijnenburg, "Ambiguity aversion and household portfolio choice puzzles: Empirical evidence," *Journal of Financial Economics*, 2016, *119* (3), 559–577.

Ellison, Glenn, "Bounded rationality in industrial organization," *Econometric Society Monographs*, 2006, *42*, 142.

Esponda, Ignacio, Emanuel Vespa, and Sevgi Yuksel, "Mental Models and Learning: The Case of Base-Rate Neglect," *American Economic Review*, 2024, *114* (3), 752–782.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

Fe, Eduardo, David Gill, and Victoria Prowse, "Cognitive skills, strategic sophistication, and life outcomes," *Journal of Political Economy*, 2022, *130* (10), 2643–2704.

Fenizia, Alessandra, "Managers and productivity in the public sector," *Econometrica*, 2022, *90* (3), 1063–1084.

Fish, Sara, Yannai A Gonczarowski, and Ran I Shorrer, "Algorithmic collusion by large language models," *arXiv preprint arXiv:2404.00806*, 2024, *7.*

Gill, David and Victoria Prowse, "Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis," *Journal of Political Economy*, 2016, *124* (6), 1619–1676.

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with measurement error: Techniques with applications to the caltech cohort study," *Journal of Political Economy*, 2019, *127* (4), 1826–1863.

Goldfarb, Avi and Mo Xiao, "Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets," *American Economic Review*, 2011, *101* (7), 3130–3161.

_ and _ , "Transitory shocks, limited attention, and a firm's decision to exit," Technical Report, Working paper 2016.

Green, Edward J and Robert H Porter, "Noncooperative collusion under imperfect price information," *Econometrica: Journal of the Econometric Society*, 1984, pp. 87–100.

Hastings, Justine S, "Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in Southern California," *American Economic Review*, 2004, *94* (1), 317–328.

Heckman, James J, Jora Stixrud, and Sergio Urzua, "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior," *Journal of Labor economics*, 2006, *24* (3), 411–482.

Heidhues, Paul and Botond Kőszegi, "Behavioral industrial organization," *Handbook of Behavioral Economics: Applications and Foundations 1*, 2018, *1*, 517–612.

Hoffman, Mitchell and Steven Tadelis, "People management skills, employee attrition, and manager rewards: An empirical analysis," *Journal of Political Economy*, 2021, *129* (1), 243–285.

Hortaçsu, Ali and Steven L Puller, "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market," *The RAND Journal of Economics*, 2008, *39* (1), 86–114.

Houde, Jean-François, "Spatial differentiation and vertical mergers in retail markets for gasoline," *American Economic Review*, 2012, *102* (5), 2147–2182.

Huffman, David, Collin Raymond, and Julia Shvets, "Persistent overconfidence and biased memory: Evidence from managers," *American Economic Review*, 2022, *112* (10), 3141–3175.

Ichniowski, Casey, Kathryn Shaw, Giovanna Prennushi et al., "The effects of human resource practices on manufacturing performance: A study of steel finishing lines," *American Economic Review*, 1997, *87* (3), 291–313.

Jacob, Brian A and Lars Lefgren, "Can principals identify effective teachers? Evidence on subjective performance evaluation in education," *Journal of labor Economics*, 2008, *26* (1), 101–136.

Kane, Thomas J, Jonah E Rockoff, and Douglas O Staiger, "What does certification tell us about teacher effectiveness? Evidence from New York City," *Economics of Education review*, 2008, *27* (6), 615–631.

Kendall, Chad W and Constantin Charles, "Causal narratives," Technical Report, National Bureau of Economic Research 2022.

Knittel, Christopher R and Victor Stango, "Price ceilings as focal points for tacit collusion: Evidence from credit cards," *American Economic Review*, 2003, *93* (5), 1703–1729.

Lewis, Jeffrey B and Drew A Linzer, "Estimating regression models in which the dependent variable is based on estimates," *Political analysis*, 2005, *13* (4), 345–364.

List, John A, "Does market experience eliminate market anomalies?," *The Quarterly Journal of Economics*, 2003, *118* (1), 41–71.

_ **and Charles F Mason**, "Are CEOs expected utility maximizers?," *Journal of Econometrics*, 2011, *162* (1), 114–123.

Luco, Fernando, "Who benefits from information disclosure? the case of retail gasoline," *American Economic Journal: Microeconomics*, 2019, *11* (2), 277–305.

Malmendier, Ulrike and Geoffrey Tate, "Behavioral CEOs: The role of managerial overconfidence," *The Journal of Economic Perspectives*, 2015, *29* (4), 37–60.

McKenzie, David and Christopher Woodruff, "Business practices in small firms in developing countries," *Management Science*, 2017, *63* (9), 2967–2981.

Metcalfe, Robert D, Alexandre B Sollaci, and Chad Syverson, "Managers and Productivity in Retail," Technical Report, National Bureau of Economic Research 2023.

Minni, Virginia, "Making the invisible hand visible: Managers and the allocation of workers to jobs," Technical Report, Centre for Economic Performance, LSE 2023.

Nagel, Rosemarie, "Unraveling in guessing games: An experimental study," *The American Economic Review*, 1995, *85* (5), 1313–1326.

Noel, Michael D, “Edgeworth price cycles: Evidence from the Toronto retail gasoline market,” *The Journal of Industrial Economics*, 2007, *55* (1), 69–92.

Oprea, Ryan, “What makes a rule complex?,” *American Economic Review*, 2020, *110* (12), 3913–3951.

Pai, Mallesh and Karsten Hansen, “Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms,” Technical Report, CEPR Discussion Papers 2020.

Proto, Eugenio, Aldo Rustichini, and Andis Sofianos, “Intelligence, Errors, and Cooperation in Repeated Interactions,” *The Review of Economic Studies*, 2022, *89* (5), 2723–2767.

Schmidt, Frank L and John Hunter, “General mental ability in the world of work: occupational attainment and job performance.,” *Journal of Personality and Social Psychology*, 2004, *86* (1), 162.

Slade, Margaret E, “Vancouver’s gasoline-price wars: An empirical exercise in uncovering supergame strategies,” *The Review of Economic Studies*, 1992, *59* (2), 257–276.

Stango, Victor and Jonathan Zinman, “Behavioral biases are temporally stable,” Technical Report, National Bureau of Economic Research 2020.

Strulov-Shlain, Avner, “More than a penny’s worth: Left-digit bias and firm pricing,” *Review of Economic Studies*, 2023, *90* (5), 2612–2645.

Tadelis, Steven, Christopher Hooton, Utsav Manjeer, Daniel Deisenroth, Nils Wernerfelt, Nick Dadson, and Lindsay Greenbaum, “Learning, Sophistication, and the Returns to Advertising: Implications for Differences in Firm Performance,” Technical Report, National Bureau of Economic Research 2023.

TANGNEY, June P, Roy F BAUMEISTER, and Angie Luzio BOONE, “High self-control predicts good adjustment, less pathology, better grades, and interpersonal success,” *Journal of Personality*, 2004, *72* (2), 271–324.

Online Appendix

Table of Contents

A Measuring manager traits

A.1 Details on variables and survey waves

Table 5: Variables by data source

Category	Variable	Wave 1	Wave 2	Wave 3	Admin
Cognitive Measures	Raven's matrices	✓	✓	✓	
	Numeracy assessment	✓	✓	✓	
Mental Models	money request game	✓	✓	✓	
	Believed influence	✓			
	Price-matching	✓			
	Profit narratives		✓		
	Competitor price response			✓	
Manager Characteristics	Noncognitive skills	✓	✓		
	Gender, age	✓	✓	✓	
	Experience	✓	✓	✓	
Self-reported Pricing	Preference for lower prices		✓		
	Frequency of price cut requests		✓		
Performance	Profits (fuel and non-fuel)				✓
	Prices				✓
	Price wars				✓
Station Characteristics	Location type				✓
	Number of competitors	✓			
	Market share	✓			
	Station area				✓
	Open 24 hours	✓			
	Station rented by partner company				✓

Notes: This table reports the availability of variables across different data sources.

A.2 Raven test example

Figure 12: An example of the Raven progressive Matrices Test

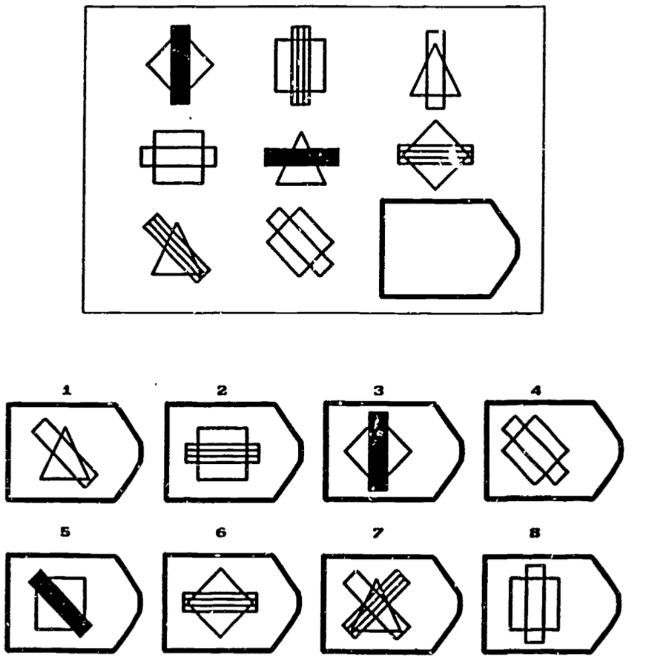

Notes: This figure shows a sample problem from Raven's Progressive Matrices Test (Source: Carpenter et al., 1990). The task requires identifying patterns across rows and columns to determine which of the eight options (numbered 1-8) correctly completes the 3×3 matrix. This example is governed by three rules: (A) each row contains three geometric figures (diamond, triangle, square) distributed across entries; (B) each row contains three textured lines (dark, striped, clear) distributed across entries; and (C) line orientation is constant within rows but varies between rows (vertical, horizontal, oblique). The correct answer is option 5. Participants who identify only some rules may select incorrect options. For instance, identifying only rule A might lead to choosing option 2 or 8, while identifying only rules B might lead to option 3. More difficult Raven tests involve additional rules or more complex pattern relationships.

A.3 Details on factor analysis of manager traits

This section provides more details on the factor analysis underlying our construction of cognitive and noncognitive skills measures.

An initial factor analysis pooling all manager traits shows that the cognitive ability measures mainly load on one factor, while other manager traits mainly load on other factors, which motivates the distinction between cognitive skills and noncognitive skills. Specifically, Table A.1 shows the three most important factors for our first survey wave. The light green shading in the table shows that most of the items from the cognitive ability measures have the highest loadings on the second factor, while the other manager traits have their highest loadings on the first or third factor. Table A.2 shows a similar structure in the second survey wave, although in this case some of the emotional intelligence items load on the second factor along with cognitive ability (we check robustness to including emotional intelligence in our cognitive skills measure, as noted in Section 2.).

Turning to analysis of cognitive skills, in both the first and second survey waves we find that there is one main factor with eigenvalue greater than 1 underlying the two cognitive ability measures (Raven test and the numeracy question). Table A.3 presents the factors for the two survey waves, and shows robustness in the sense that loadings of individual items are relatively similar across waves, e.g, items 4 and 6 on the Raven test have the lowest loadings in both waves. For each survey wave we use the predicted value of the corresponding factor for each manager as our cognitive skills measure (manager responses to the individual items are weighted by the loadings to create a score).

For both survey waves, a factor analysis of other manager traits reveals one main factor, with eigenvalue well above 1, and a second factor with eigenvalue equal to or slightly below 1 (Table A.4). The shading in the table shows that individual measures and items load on the first factor in a similar way across the two waves. We use the predicted value of the first factor from each wave to construct our corresponding measure of a manager's noncognitive skills.

Table A.1: First three factors of manager traits, first survey wave (largest loading in green)

Variable	**Factor 1**	**Factor 2**	**Factor 3**
Raven 1	0.3295	0.3624	0.0912
Raven 2	0.2516	0.2211	0.0384
Raven 3	0.4002	0.3709	0.0591
Raven 4	0.1189	0.2217	0.0879
Raven 5	0.2768	0.3582	0.1015
Raven 6	0.0307	0.0691	0.0533
Raven 7	0.3905	0.4350	0.1017
Raven 8	0.4300	0.5038	0.1082
Raven 9	0.3320	0.3701	0.0816
Coin toss	0.1204	0.1540	-0.0179
Risk taking	-0.0043	-0.1825	0.0849
Patience	0.1433	-0.0465	0.0131
Trust	0.4019	-0.0779	-0.0343
Negative reciprocity	-0.1735	0.1142	0.3652
Positive reciprocity	0.4727	0.0513	-0.1174
Altruism	0.1508	-0.2062	-0.0187
Extraversion	0.3486	-0.4169	0.2913
Agreeableness	0.4734	-0.2005	-0.0452
Conscientiousness	0.5296	-0.3763	0.0612
Neuroticism	-0.4097	0.3218	-0.0600
Openness to experience	0.1073	-0.2974	0.3715
Competitiveness	0.2168	-0.2029	0.3058
Confidence	0.4662	-0.2709	0.2133
Locus of control	0.4802	-0.1730	-0.1419
Authority preference	-0.2217	0.0200	0.3132
Self-control	0.5932	-0.2654	-0.2697
Procrastination	-0.3231	0.0883	0.2931
Ambiguity aversion	0.0026	0.0478	0.0543
Emotional IQ 1	0.2119	0.1210	-0.0449
Emotional IQ 2	0.0598	0.0745	0.0061
Emotional IQ 3	0.0808	0.0446	0.0013
Emotional IQ 4	0.1773	0.1508	-0.0321
Emotional IQ 5	0.1770	0.1293	-0.0200
Emotional IQ 6	0.2389	0.1745	-0.0436
Emotional IQ 7	0.0411	0.0225	-0.0249
Emotional IQ 8	0.1271	0.0460	-0.0127
Number of managers	13,655	13,655	13,655
Eigenvalue	3.34	2.07	0.88

Notes: The table reports the corresponding factor loadings for the first three factors from a factor analysis of all manager characteristics, including both cognitive and noncognitive characteristics, in the first survey wave. For each characteristic, the factor on which it has the highest loading has the loading color-coded green. Note that only the first two factors in the analysis have eigenvalues greater than 1.

Table A.2: First three factors of manager traits, second survey wave (largest loading in green)

Variable	Factor1	Factor2	Factor3
Raven 1	0.2439	0.4017	0.0849
Raven 2	0.2455	0.3356	0.0543
Raven 3	0.2986	0.3822	-0.0004
Raven 4	0.1201	0.3061	0.1451
Raven 5	0.2577	0.4325	0.1218
Raven 6	0.0574	0.1392	0.1033
Raven 7	0.3162	0.5017	0.0839
Raven 8	0.3102	0.5212	0.1063
Raven 9	0.1536	0.2946	0.1195
Coin toss	0.1045	0.1600	-0.0342
Risk taking	0.0942	-0.0806	0.0775
Patience	0.1632	-0.0554	0.0007
Trust	0.3489	-0.0502	-0.1236
Negative reciprocity	-0.2718	0.0994	0.3453
Positive reciprocity	0.4131	0.1088	-0.1627
Altruism	0.1931	-0.1816	-0.0524
Extraversion	0.3992	-0.3826	0.3162
Agreeableness	0.5128	-0.1526	-0.0280
Conscientiousness	0.5580	-0.2982	0.0646
Neuroticism	-0.4522	0.2813	-0.1485
Openness to experience	0.2445	-0.2828	0.3719
Competitiveness	0.2362	-0.1560	0.2704
Confidence	0.4796	-0.2002	0.1594
Locus of control	0.5117	-0.1005	-0.1407
Authority preference	-0.2153	-0.0468	0.2676
Self-control	0.6505	-0.1708	-0.2209
Procrastination	-0.3935	0.0606	0.2839
Ambiguity aversion	0.0614	0.0740	0.0184
Emotional IQ 1	0.1648	0.1273	-0.0406
Emotional IQ 2	0.0402	0.0879	-0.0005
Emotional IQ 3	0.0581	0.0734	-0.0012
Emotional IQ 4	0.1394	0.1858	-0.0021
Emotional IQ 5	0.1422	0.1565	-0.0135
Emotional IQ 6	0.1709	0.2068	-0.0418
Emotional IQ 7	0.0270	-0.0143	-0.0113
Emotional IQ 8	0.0799	0.0595	-0.0258
Number of managers	15,872	15,872	15,872
Eigenvalue	3.24	2.09	0.85

Notes: The table reports the corresponding factor loadings for the first three factors from a factor analysis of all manager characteristics, including both cognitive and noncognitive characteristics, in the second survey wave. For each characteristic, the factor on which it has the highest loading has the loading color-coded green. Note that only the first two factors in the analysis have eigenvalues greater than 1.

Table A.3: First cognitive skills factor, first and second survey waves

Variable	**First wave**	**Second wave**
Raven 1	0.4911	0.4795
Raven 2	0.3290	0.4185
Raven 3	0.5273	0.4630
Raven 4	0.2660	0.3518
Raven 5	0.4725	0.5247
Raven 6	0.0806	0.1647
Raven 7	0.6044	0.6047
Raven 8	0.6944	0.6268
Raven 9	0.5218	0.3522
Coin toss	0.1814	0.1751
Number of managers	13,655	15,872
Eigenvalue	2.08	1.96

Notes: The table reports the corresponding factor loadings for the primary cognitive skills factor in the first and second survey waves. The factor analysis used indicators for correctly answering each of the 9 items of the Raven test and the question about the probability of heads on a coin toss. For each wave, there was only one factor with an eigenvalue greater than 1.

Table A.4: First two factors of noncognitive skills, first and second survey waves

Variable	**First factor** First wave	Second wave	**Second factor** First wave	Second wave
Risk taking	0.0718	0.1145	0.1664	0.0870
Patience	0.1503	0.1726	-0.0084	0.0015
Trust	0.4023	0.3488	-0.1213	-0.1373
Negative reciprocity	-0.2151	-0.2972	0.2274	0.2669
Positive reciprocity	0.4064	0.3536	-0.2502	-0.2574
Altruism	0.2243	0.2439	0.0717	0.0266
Extraversion	0.4816	0.4964	0.3943	0.3901
Agreeableness	0.5148	0.5335	-0.0383	-0.0237
Conscientiousness	0.6372	0.6238	0.1369	0.1200
Neuroticism	-0.5053	-0.5162	-0.1418	-0.2002
Openness to experience	0.2083	0.3163	0.4516	0.4079
Competitiveness	0.2724	0.2683	0.2905	0.2601
Confidence	0.5318	0.5158	0.1883	0.1502
Locus of control	0.5124	0.5183	-0.1207	-0.1397
Autonomy preference	-0.2169	-0.1922	0.2603	0.2672
Self-control	0.6580	0.6752	-0.1852	-0.1944
Procrastination	-0.3351	-0.3975	0.2211	0.2608
Ambiguity aversion	-0.0220	0.0313	0.0203	-0.0323
Emotional IQ 1	0.1420	0.1144	-0.1659	-0.1428
Emotional IQ 2	0.0249	0.0106	-0.0742	-0.0751
Emotional IQ 3	0.0561	0.0323	-0.0621	-0.0707
Emotional IQ 4	0.1039	0.0764	-0.2088	-0.1968
Emotional IQ 5	0.1129	0.0873	-0.1919	-0.1825
Emotional IQ 6	0.1443	0.0936	-0.2143	-0.1912
Emotional IQ 7	0.0314	0.0322	-0.0642	-0.0229
Emotional IQ 8	0.0995	0.0581	-0.0990	-0.0906
Number of managers	13,655	15,872	13,655	15,872
Eigenvalue	2.94	3.02	1.02	0.97

Notes: The table reports the corresponding factor loadings for the first and second noncognitive skills factors in the first and second survey waves. Columns (2) and (3) show the first factors for the two survey waves, while Columns (4) and (5) show the second factors. The factor analysis for each wave used answers to each of the noncognitive characteristic measures, as well as indicators for correctly answering each of the 8 emotional intelligence questions. For each wave, there were only two noncognitive skills factors with eigenvalues greater than or approximately equal to 1. Positive loadings are color coded in blue, negative in red.

A.4 Measurement error in cognitive and noncognitive skills

The repetition of the measures of manager traits across two waves offers the possibility of assessing and correcting for measurement error. A caveat is that, due to company policy regarding data privacy, we do not have personal identifiers in the first two waves that can directly match them to each other. To overcome this constraint, we implemented a matching procedure based on self-reported demographic variables (age, tenure) and employment history (specific stations worked at during particular years). This approach allowed us to match approximately 60 percent of managers between the first and second waves (this is similar to the overlap we observe for waves two and three, for which we have more direct identifiers, so it appears the matching works well).

Quantifying measurement error in cognitive and noncognitive skills serves two purposes. First, it helps us understand the extent to which attenuation bias might cause our estimates of the relationship between cognitive skills and various outcomes to represent lower bounds of the true effects. Second, it addresses a potential concern in our regression analyses: when both cognitive and noncognitive skills are included as predictors, measurement error in noncognitive skills could bias the estimated coefficient on cognitive skills (see, e.g., Gillen et al., 2019).

Our measurement error calculations find greater error for cognitive skills than for noncognitive skills, and imply non-trivial attenuation. Specifically, we estimate that correlations between outcomes and cognitive skills are attenuated by approximately 35 percent. The same calculation for noncognitive skills shows that the observed correlations are attenuated by about 17 percent.

For regression analyses that include both cognitive and noncognitive skills as independent variables, we can check robustness to correcting for measurement error using the Obviously Related Instrumental Variables (ORIV) approach, instrumenting for manager traits measured in one survey wave with the same traits measured in the other survey wave (Gillen et al., 2019; see also Stango and Zinman, 2020). We perform this correction for a wide range of our main regression analyses, which use cognitive skills and other manager traits as independent variables: (1) Impact of cognitive skills on having a narratives that mentions high price (Figure A.1); (2) impact of cognitive skills on self-reported desire to cut prices (Figure A.2); (3) impact of cognitive skills on self-reported frequency of proposing price cuts (Figure A.3); (4) impact of cognitive skills on actual monthly price ratio (Figure A.4); (5) impact of cognitive skill on SDID treatment effects on monthly price ratio from our event study analysis (Figure A.5). We find no evidence that the coefficient on cognitive skills is reduced in any of these. Instead, it is always larger with ORIV, ususally substantially so, consistent with the main bias of the cognitive skills coefficient being attenuation.

Figure A.1: High price narrative as a function of cognitive skills: ORIV approach

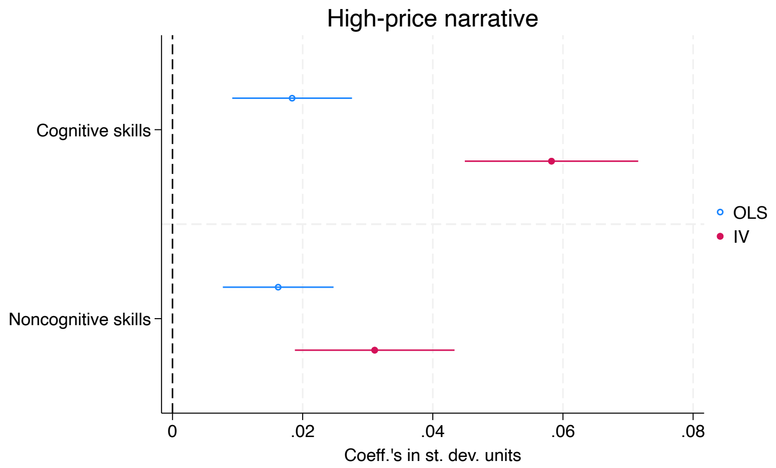

Notes: This figure reports coefficients from regressions examining the relationship between cognitive skills and mentioning the high price narrative, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable equals to 1 if the manager mentioned high price as the cause to high fuel profit. The sample consists of managers who participated in both survey waves and could be matched through demographic identifiers. The "OLS" estimates show results from standard OLS regression on this matched sample. The "IV" estimates show results using the Obviously Related Instrumental Variables (ORIV) approach from Gillen et al. (2019), which corrects for measurement error by instrumenting each skill measure from one survey wave with the corresponding measure from the other survey wave. Controls include other manager traits (experience, gender, and age), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators), and interacted month and district fixed effects.

Figure A.2: Desire to cut prices as a function of cognitive skills: ORIV approach

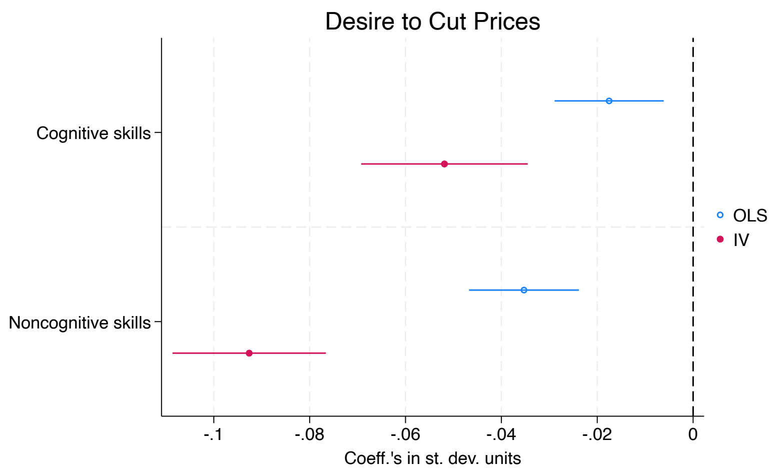

Notes: This figure reports coefficients from regressions examining the relationship between manager skills and desire to cut prices, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the self-reported desire to charge lower than the default price set by upper-level management. For the rest of the details, see Figure A.1.

Figure A.3: Frequency of proposing price cuts as a function of cognitive skills: ORIV approach

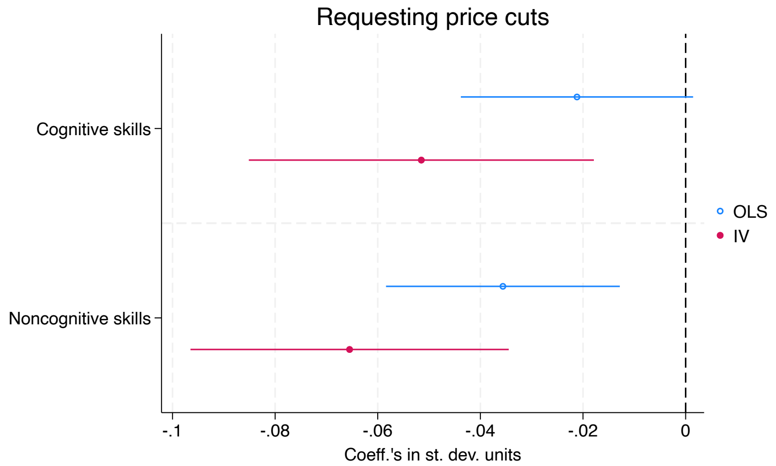

Notes: This figure reports coefficients from regressions examining the relationship between manager skills and pricing behavior, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the frequency to propose a price cut to upper-level management. For the rest of the details, see Figure A.1.

Figure A.4: Pricing behavior as a function of cognitive skills: ORIV approach

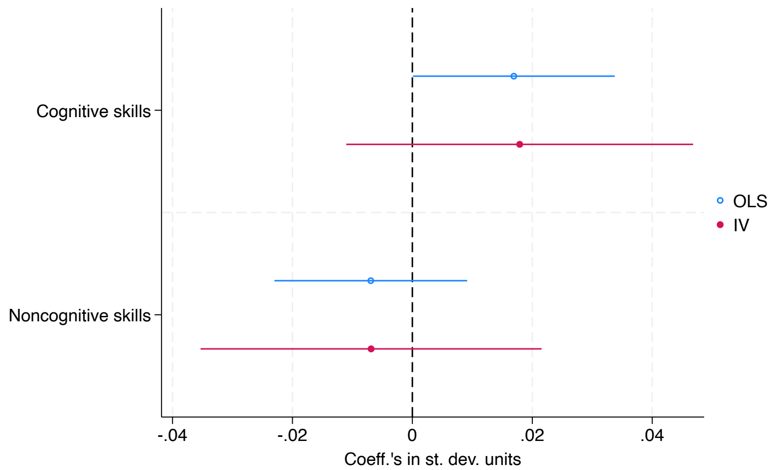

Notes: This figure reports coefficients from regressions examining the relationship between manager skills and pricing behavior, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the standardized monthly price ratio. For the rest of the details, see Figure A.1.

Figure A.5: SDID treatment effects as a function of cognitive skills: ORIV approach

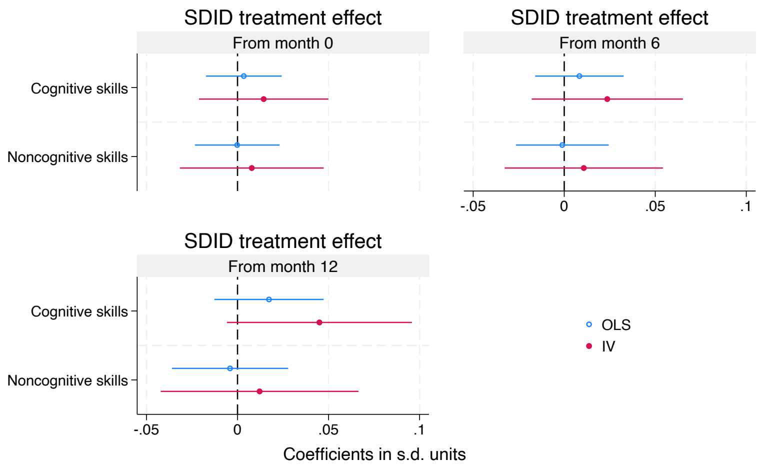

Notes: This figure reports coefficients from regressions examining the relationship between manager cognitive skills and SDID treatment effects, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the synthetic difference-in-differences (SDID) treatment effect of the new manager on the price ratio. For details of the SDID method, see Section C.2.1. For the rest of the details, see Figure A.1.

B Mental models

B.1 Robustness of results on the money request game and cognitive skills

The results from the money request game, and the link to cognitive skills, are similar across survey waves. As in the first survey wave, we find in both the second and third waves that the modal request is $5, the optimal request is $4, and requesting $4 is strongly predicted by cognitive skills (Figures B.1 and B.2). The overlap between survey waves is about 60 percent between adjacent waves, so the results come from substantially different samples each time.

Figure B.1: Behavior in the money request game and cognitive skills (second wave)

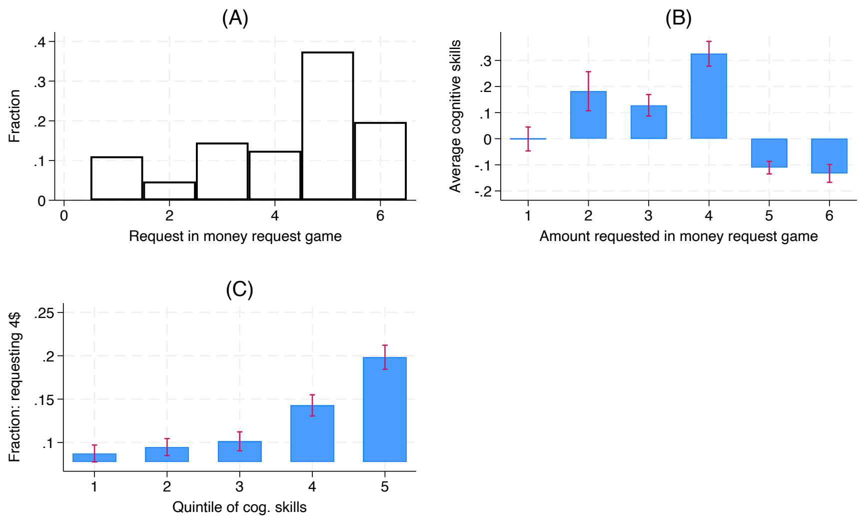

Notes: This table reports the behaviors in the money request game from the second survey wave. See Figure 1 for details.

Figure B.2: Behavior in the money request game and cognitive skills (third wave)

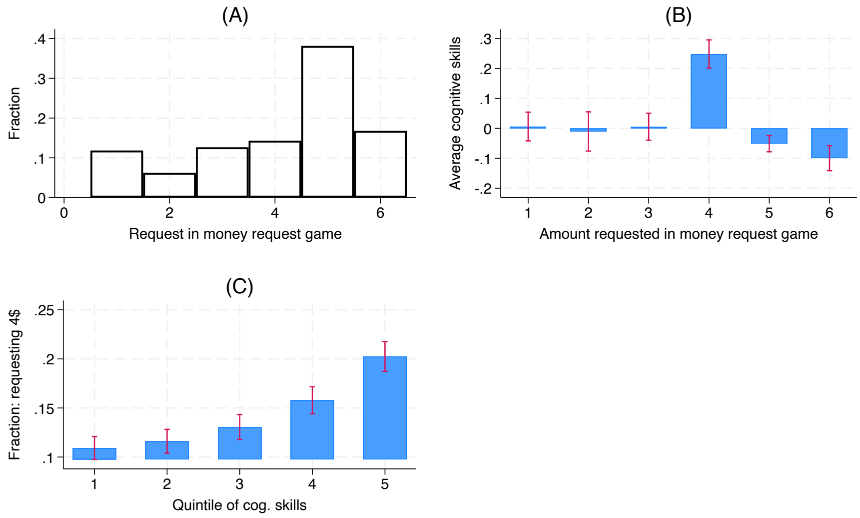

Notes: This table reports the behaviors in the money request game from the third survey wave. See Figure 1 for details.

B.2 Relationships of cognitive skills to mental models across market conditions

Figure B.3: Relationship of mental models to cognitive skills, by number of competitors

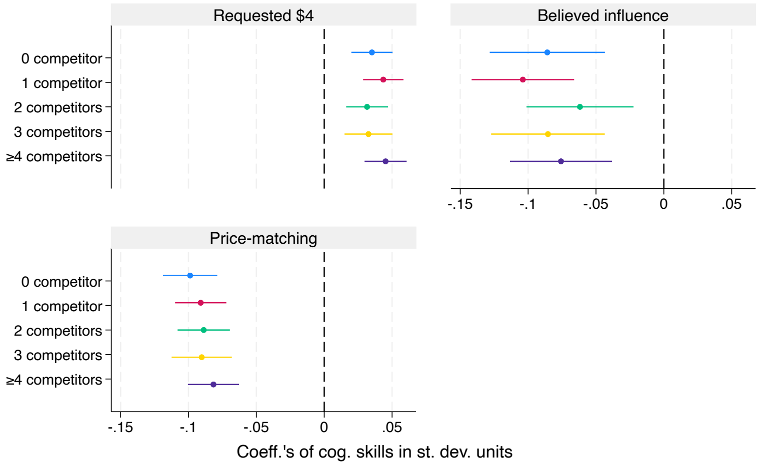

Notes: This figure reports coefficients from OLS regressions on three mental model measures by number of competitors, with 95% confidence intervals. The three dependent variables are: requesting $4 in the money request game, believed influence over fuel sales, and believed optimality of price-matching strategy. Coefficients are for cognitive skills as an explanatory variable. Each regression controls for other traits of managers (noncognitive skills, age, gender, experience), and station and location characteristics (open 24 hours, whether the company rents the station, thirteen location type indicators, district fixed effects). The groupings of competitors are based on quintiles of the distribution of number of competitors.

Figure B.4: Relationship of mental models to cognitive skills, by location types

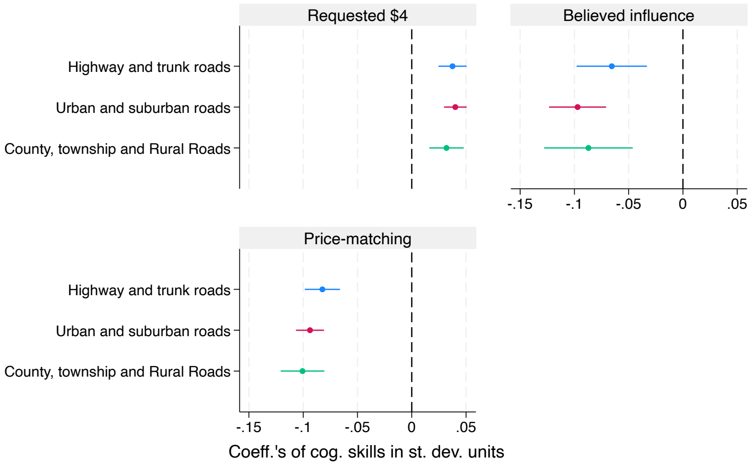

Notes: This figure reports coefficients from OLS regressions on three mental model measures by location types, with 95% confidence intervals. The three dependent variables are: requesting $4 in the money request game, believed influence over fuel sales, and believed optimality of price-matching strategy. Coefficients are for cognitive skills as an explanatory variable. Each regression controls for other traits of managers (noncognitive skills, age, gender, experience), and station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, district fixed effects).

Figure B.5: Probability of mentioning high price narrative and cognitive skills, by market conditions and location type

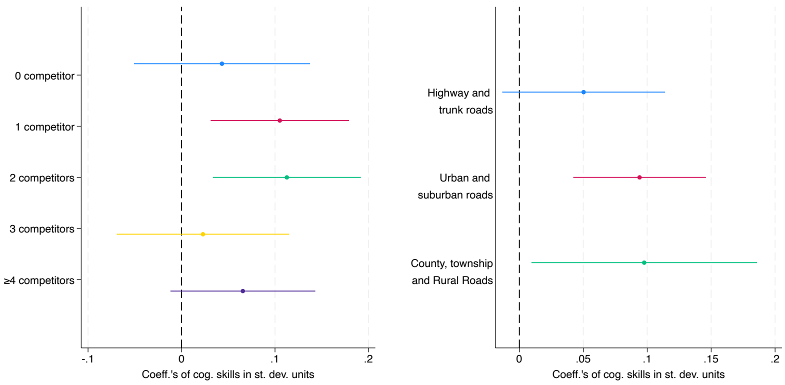

Notes: This figure reports marginal effects from Probit regressions, with 95% confidence intervals. The dependent variable is an indicator for whether a manager mentioned the high price narrative. Coefficients are for cognitive skills as an explanatory variable. The left panel shows the coefficients by the number of competitors and the right panel shows the coefficients by location types. The groupings of competitors are based on quintiles of the distribution of number of competitors. Each regression controls for other traits of managers (noncognitive skills, age, gender, experience), and station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators).

B.3 Classification and analysis of manager narratives about causes of fuel profits

B.3.1 Details on narratives classification

In a first stage, the researchers looked at a randomly drawn sub-sample of 3,000 responses out of the total sample of more than 15,000. We accumulated a list of distinct categories of causes of high fuel profits mentioned by the managers, e.g., keeping fuel prices high to maintain high margins, charging low prices to increase sales volume, or a manager putting in a lot of effort. We found that many of these causes could be conceptualized within the relationship $\pi = (p - c) \times q$, namely as being related to either the profit margin component of profit or to the sales volume component.

Following Andre et al. (2023) we used RAs to classify the text responses. We provided the undergraduate RAs with the fifteen categories we had identified, along with examples of text belonging to each and a list of common keywords associated with each category, to serve as a rubric for classifying responses (see Appendix H for the rubric). The RAs then categorized all 15,000 responses, with two RAs looking at each item of text. The agreement rate between RAs was about 75 percent. For the results reported in the main text, conflicts in RA categorization were reconciled by the researchers.

Two categories of our rubric mentioned the importance of keeping prices high, or avoiding across-the-board discounts, in slightly different ways: (1) keeping listed price of fuel high (mentioned by 13% of managers), (2) reducing price only for select groups of customers (mentioned by 2.5%). Due to the conceptual similarity, we combined these into a single high price category for the purpose of our analysis. Notably, managers who mention (2)—giving a more nuanced explanation that across-the-board price cuts are more problematic than targeted ones—have even higher cognitive ability than those in group (1), i.e., those who mention high price in some other way.

B.3.2 Robustness: Relationship of narratives to cognitive skills

We have performed various robustness checks on the classification of narratives, which eliminate the role of researchers in classification. Results are very similar based on having a third RA, rather than researchers, reconcile the 25% of narratives for which there was disagreements between the two initial RA's (see Figure B.6). Results are also very similar if we only use the 75% of narratives for which the two RAs agreed on the classification (see Figure B.7).

We also check robustness of the link between mentioning high price and cognitive skills by varying which other causes are included in the analysis. The first regression in Figure B.8 ("All") uses the full sample and is the same as the first coefficient in Figure 5 in the text, while the second, "Frequent," matches the sample used for Panel (C) of Figure 4 in the main text. In both samples we see that cognitive skills strongly predict mentioning high price. The "Frequent without 'don't know'" addresses a potential confound: that managers who answer "don't know" might be generally less willing to exert effort, both in answering the narrative question and in completing the Raven's matrices. In this case, a greater tendency to mention "don't know" rather than high price might not indicate an effect of bounded rationality but rather lack of effort. We see, however, that even if we exclude managers who say "don't know" and focus only on those who give a specific cause, cognitive skills strongly predict mentioning high price. The "Sales vol." specification further restricts the sample, showing that among managers who mention either high price or a sales volume cause like low price, cognitive skills strongly predict whether a manager mentions high price. Finally, the "Low price" specification only includes managers who either mention high price or explicitly mention low price, and shows that cognitive skills are a significant predictor of mentioning the high price cause instead of low price.

We also examine whether our results are driven by high-skill managers mentioning more causes overall, which could mechanically make it more likely that they mention high price (number of narratives mentioned is modestly, positively correlated with cognitive skills; Spearman $\rho = 0.14$, $p < 0.001$). Figure B.9 restricts the sample to managers who mention exactly one narrative. In this restricted sample, we observe the same pattern: higher cognitive skills significantly predict mentioning the high price narrative (Regression 1), and this relationship is partially mediated by measures of mental models of competitors (Regression 2).

We also investigate whether the relationship between cognitive skills and narratives varies across different market conditions. A natural concern is that managers with different cognitive skills might have different ideas about how to achieve fuel profits because they face different market conditions. Appendix Figure B.5 shows, however, that the divergence of views about optimal strategies by cognitive skills is strikingly similar across a wide range of market structures; whether there are many competitors, or only a few, or whether the station is on a highway or in a small town, higher-skill managers have a tendency to favor keeping prices high. This suggests that, instead, there is a general tendency for cognitive skills to matter for mental models of profits, potentially working through different underlying mental models of how competitors behave.

Finally, Figure B.10 investigates whether the results are robust to using trait measures from the second survey wave, rather than the first wave; using traits from the first wave, which allows including beliefs about influence and optimality of price-matching, restricts the sample to managers who answer both waves. However, we observe the same significant relationship between cognitive skills and mentioning the high price narrative using the full sample of the second wave with cognitive skills, money request, and other controls measured in the second wave.

Figure B.6: Narrative measure of mental models for high fuel profits (Third RA reconciles disagreement)

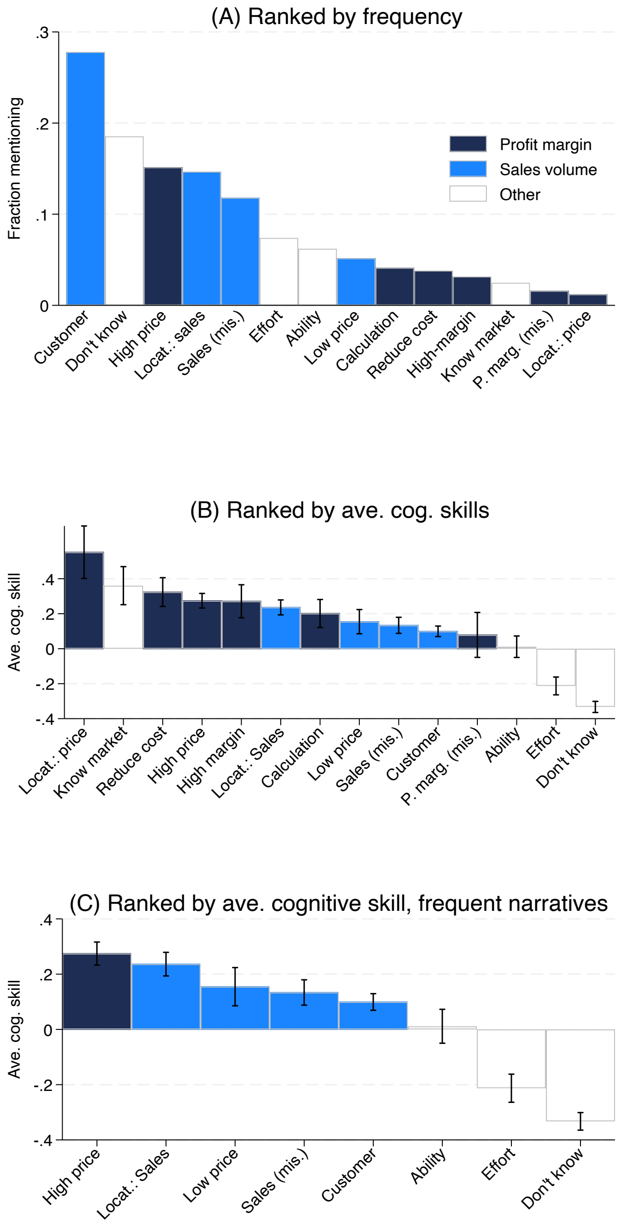

Notes: This figure reports the frequency and average manager cognitive skills of narratives about what drives high fuel profits. Here a third RA, rather than researchers, reconcile the 25% of narratives for which there was disagreements between the two initial RA's. See Figure 4 for the rest of the details.

Figure B.7: Narrative measure of mental models for high fuel profits (RA's agree)

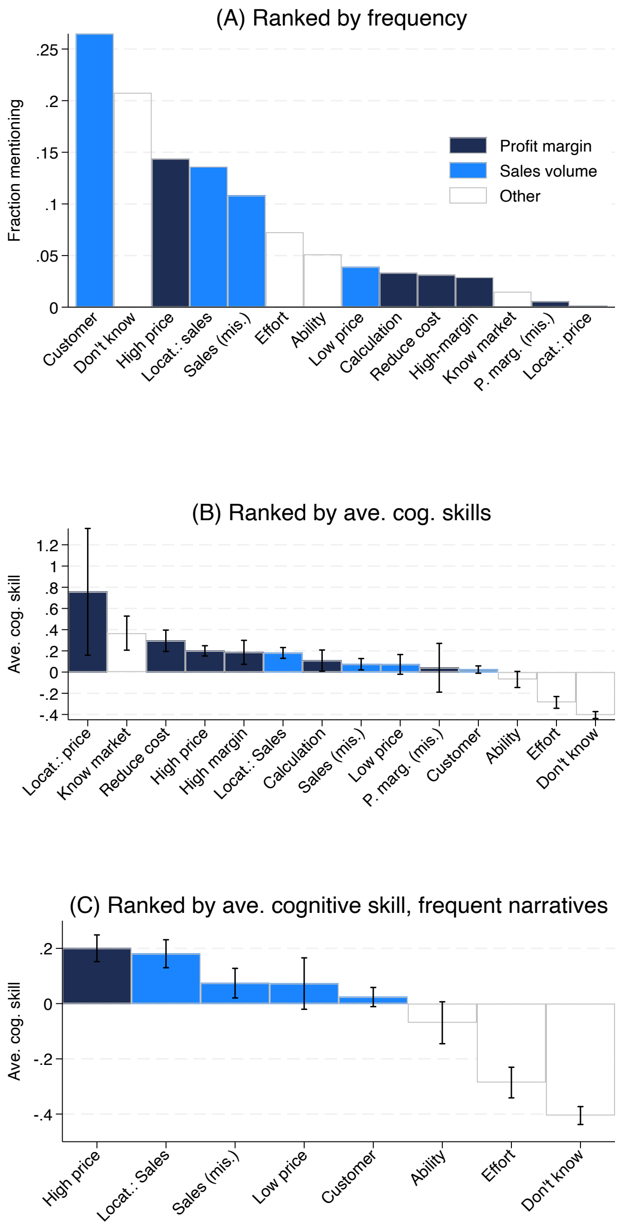

Notes: This figure reports the frequency and average manager cognitive skills of narratives about what drives high fuel profits. This figure only shows the 75% of cases with full RA agreement. See Figure 4 for the rest of the details.

Figure B.8: Probability of mentioning high price narrative and cognitive skills, by sample

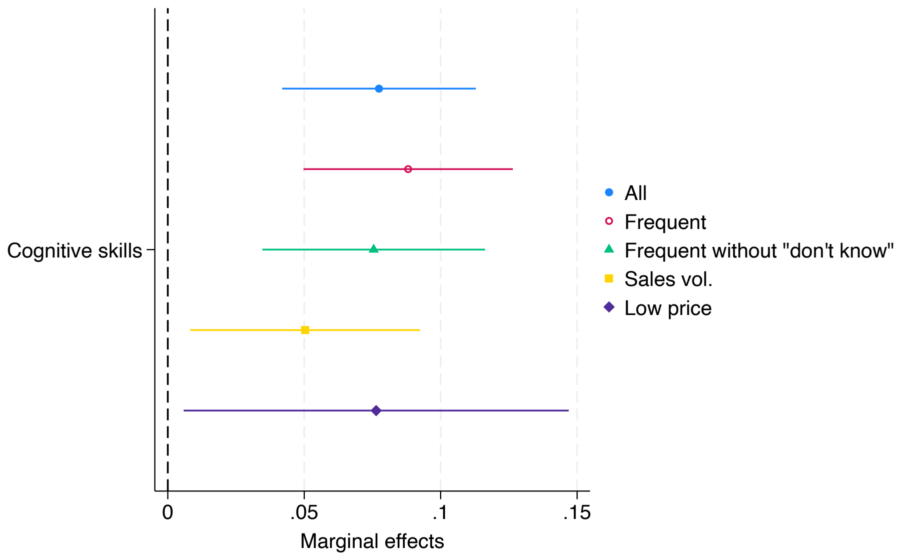

Notes: The figure reports marginal effects from Probit regressions, with 95% confidence intervals. Each coefficient is from a separate regression, and shows how a 1 s.d. increase in cognitive skills translates into the probability of mentioning the high price narrative, relative to mentioning a narrative from the respective baseline group. "All" uses all managers who mention either high price, or one alternative narrative. "Frequent" only uses managers who mention either high price or one of the relatively frequent (> 5% frequency) alternative narratives. "Frequent without 'don't know'" is the same as the "Frequent" sample except that the category of "don't know" is not included. "Sales vol." uses managers who mention either high price or narratives from the sales volume category. "Low price" only uses managers who mention either high price or the low price narrative. For the rest of the details, see Figure 5.

Figure B.9: Probability of mentioning high price narrative and cognitive skills, managers mentioning a single narrative

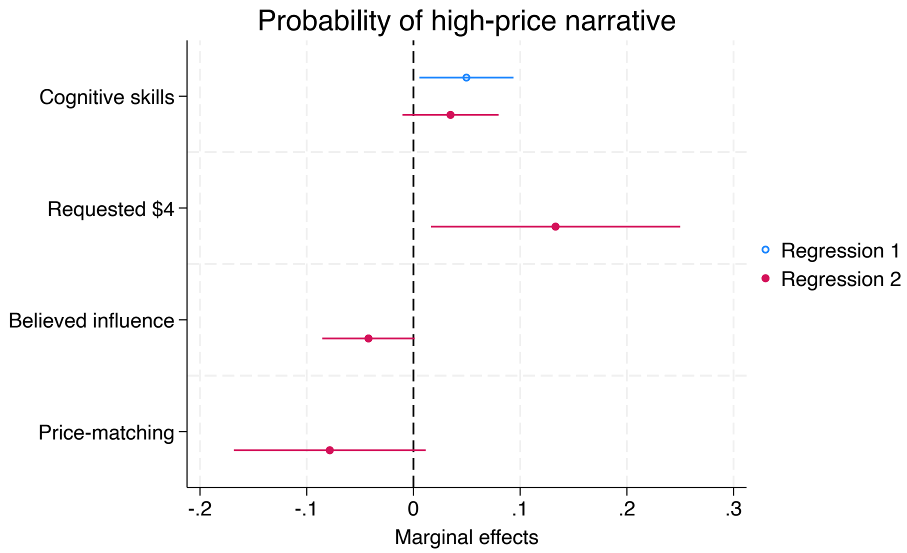

Notes: The figure reports marginal effects from Probit regressions, with 95% confidence intervals. The sample is restricted to managers who mention only a single cause. For the rest of the details, see Figure 5.

Figure B.10: Probability of mentioning high price narrative and cognitive skills, second wave only

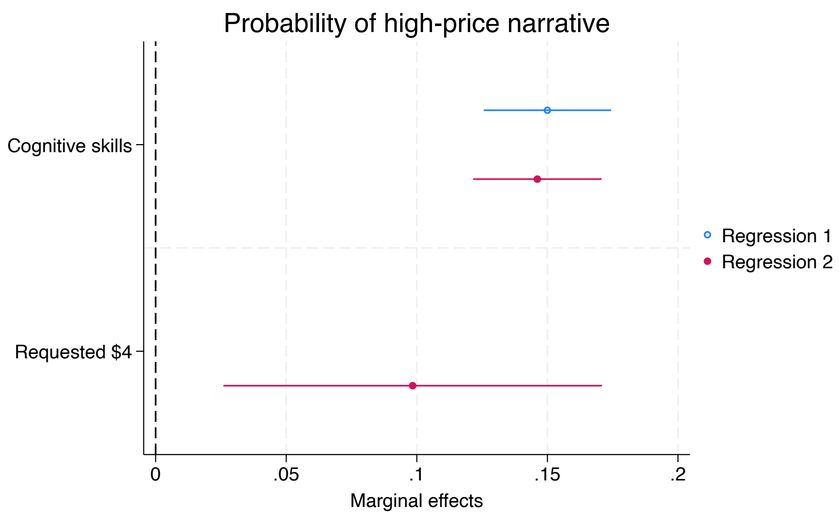

Notes: The figure reports marginal effects from Probit regressions using cognitive skills and mental models measured in the second wave, with 95% confidence intervals. The sample consists of managers participating in the second survey wave, with dependent and independent variables measured in the second wave. For the rest of the details, see Figure 5.

B.4 Additional figure and tables

Table B.1: Balance table

	Cog. skills above median (1)	Cog. skills below median (2)	Difference (1) − (2)
Station area	832.549	847.698	-15.149
	(2,235.717)	(2,343.108)	(45.393)
Open 24 hours	0.710	0.718	-0.008
	(0.454)	(0.450)	(0.008)
Number of major competitors	0.937	0.987	-0.050**
	(1.078)	(1.094)	(0.020)
Number of independent competitors	1.243	1.281	-0.038
	(1.401)	(1.483)	(0.026)
Number of nearby partner stations	0.885	0.934	-0.049**
	(1.050)	(1.059)	(0.019)
Station rented by partner company	0.184	0.167	0.017**
	(0.387)	(0.373)	(0.007)
Highway & trunk roads	0.367	0.368	-0.001
	(0.482)	(0.482)	(0.009)
Urban & suburban roads	0.447	0.436	0.011
	(0.497)	(0.496)	(0.009)
County, township & rural roads	0.187	0.196	-0.009
	(0.390)	(0.397)	(0.007)
Observations	6,029	5,961	11,990

Notes: *** p<0.01, ** p<0.05, * p<0.1.

Table B.2: Mental models and cognitive skills

VARIABLES	Requested \$4 (1)	Believed influence (2)	Price-matching (3)
Cognitive skills	0.038***	-0.082***	-0.090***
	(0.004)	(0.009)	(0.004)
Noncognitive skills	-0.001	0.144***	0.023***
	(0.003)	(0.009)	(0.004)
Experience	-0.001	0.052***	0.018***
	(0.004)	(0.012)	(0.006)
Female	-0.014*	-0.059***	0.004
	(0.007)	(0.020)	(0.010)
Age	-0.006	-0.019	-0.057***
	(0.004)	(0.012)	(0.006)
# major competitors	0.004	0.021**	0.032***
	(0.004)	(0.010)	(0.005)
# nearby partner stations	0.000	-0.013	-0.004
	(0.003)	(0.009)	(0.005)
# independent competitors	-0.001	-0.014**	-0.003
	(0.002)	(0.006)	(0.003)
Observations	11,990	11,990	11,990
R-squared	0.018	0.061	0.075

Notes: *** p<0.01, ** p<0.05, * p<0.1. Robust standard errors are in parentheses. The table reports coefficients from OLS regression on mental models. Both columns control for station location indicators, whether the station is rented by partner company, open 24 hours, and district fixed effects. The sample consists of managers participating in both the first and second survey waves.

Table B.3: High-price narrative, cognitive skills, and mental models of competitors

VARIABLES	Probability of High-price narrative				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	0.019***	0.018***	0.018***	0.016***	0.015***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Requested $4		0.026**			0.025*
		(0.013)			(0.013)
Believed influence			-0.010**		-0.009**
			(0.004)		(0.004)
Price-matching				-0.027***	-0.025***
				(0.009)	(0.009)
Noncognitive skills	0.017***	0.017***	0.018***	0.017***	0.019***
	(0.004)	(0.004)	(0.005)	(0.004)	(0.005)
Experience	0.010*	0.010*	0.010*	0.010*	0.011*
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Female	0.013	0.013	0.012	0.013	0.013
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Age	-0.006	-0.005	-0.006	-0.007	-0.007
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
# major competitors	-0.005	-0.005	-0.004	-0.004	-0.004
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
# nearby partner stations	0.001	0.001	0.001	0.001	0.001
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
# independent competitors	-0.002	-0.002	-0.002	-0.002	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Observations	7,428	7,428	7,428	7,428	7,428

Notes: *** p<0.01, ** p<0.05, * p<0.1. The table reports marginal effects from Probit regressions. The dependent variable is an indicator for whether a manager mentioned the high price narrative. The sample consists of managers participating in both the first and second survey waves, with the narrative measure elicited in the second wave, but the independent variables in the first wave. All columns controls for station location indicators, whether the station is rented by partner company, open 24 hours, and district fixed effects.

C Pricing behavior and cognitive skills

C.1 Cognitive skills and self-reported pricing behaviors

Figure C.1: Desire to cut prices, cognitive skills, and high-price narrative

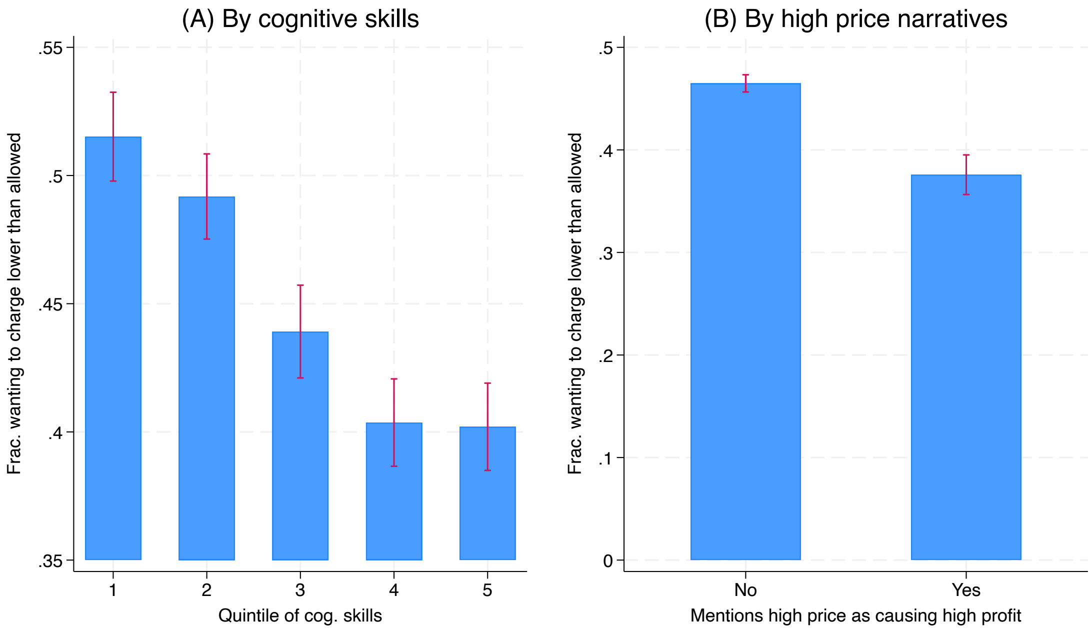

Notes: Panel (A) reports the share of managers preferring charging a lower fuel price than the default set by upper level management, categorized by quintiles of cognitive skills (quintile 5 is the highest). Panel (B) reports the fraction preferring charging a lower price by whether mentioned the high price narrative. Error bars show 95% confidence intervals.

Figure C.2: Requesting price cuts, cognitive skills, and high-price narrative

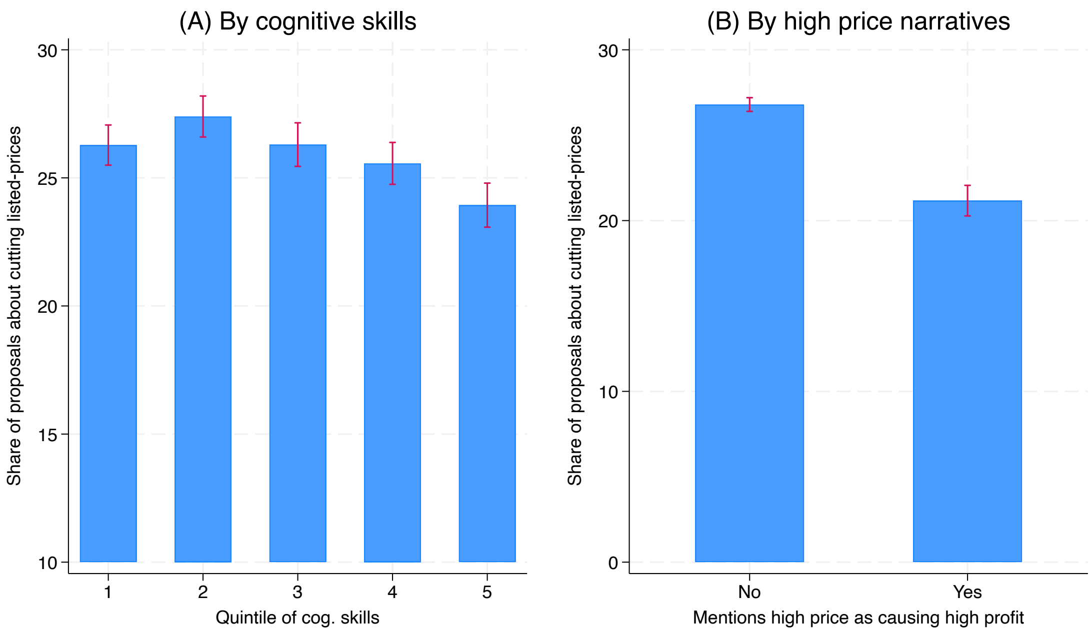

Notes: Panel (A) reports the percentages of proposals made to upper-level management about cutting listed-prices, categorized by quintiles of manager cognitive skills (quintile 5 is the best). Panel (B) reports the percentages by whether mentioned the high price narrative. Error bars show 95% confidence intervals.

Figure C.3: Probability of stating a desire to cut prices as a function of cognitive skills

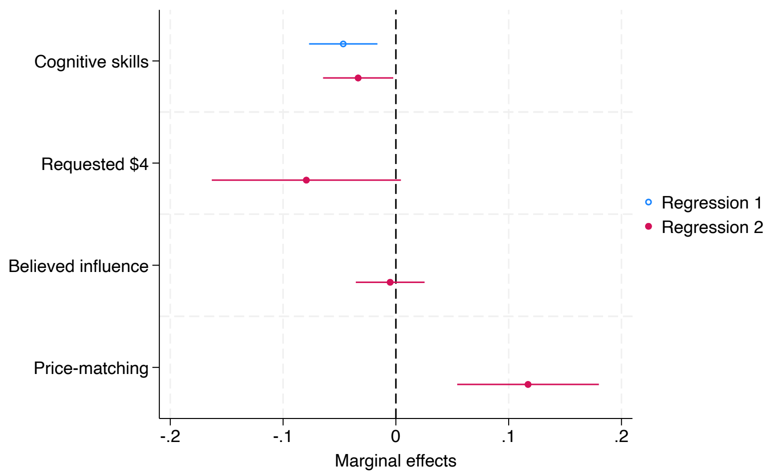

Notes: This figure reports marginal effects from Probit regressions, with 95% confidence intervals. The dependent variable equals 1 if the manager reports a tendency to cut prices relative to the default price set by upper-level management. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, experience, gender, and age), market structure (numbers and types of local competitors and location type), station characteristics (open 24 hours, and whether the station is rented by partner company), and district fixed effects. Regression 2 adds three measures of mental models of competitors.

Figure C.4: Percentage of proposals about cutting listed fuel prices as a function of cognitive skills

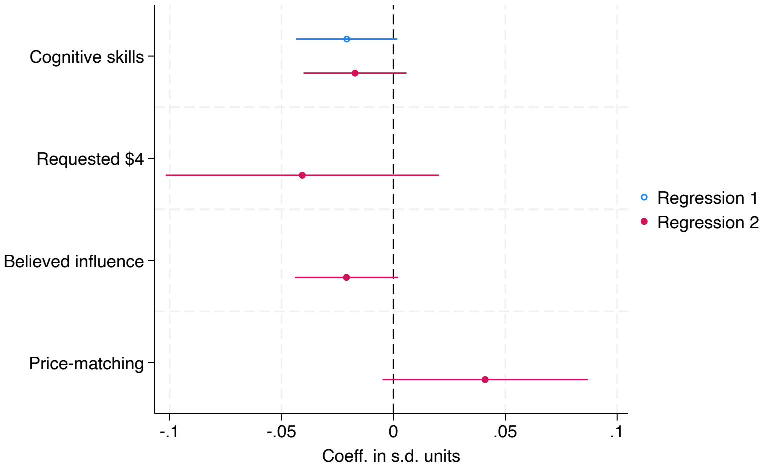

Notes: This figure reports coefficients from OLS regressions, with 95% confidence intervals. The dependent variable is the self-reported percentage of proposals that the manager makes that are to cut listed fuel prices. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, experience, gender, and age), market structure (numbers and types of local competitors and location type), station characteristics (open 24 hours, and whether the station is rented by partner company), and district fixed effects. Regression 2 adds three measures of mental models of competitors.

C.2 Additional results on actual pricing

C.2.1 Details on event study design

We use the synthetic difference-in-differences (SDID) method developed by Arkhangelsky et al. (2021) to address the concern that managers with high and low cognitive skills could be assigned to stations with systematically different unobservables, which matter for pricing. This design allows us to control for both observable and unobservable station characteristics that are reflected in pre-treatment pricing patterns.

For each manager-change event e, we let $i \in \mathscr{L}$ index stations that never gone through a manager change event between 2019 and 2022, $t \in \{-T_0, \ldots, -1, 0\}$ index months relative to the arrival of the new manager where $t = 0$ is the first post-arrival month, and P_{it} denotes the monthly price level.The SDID treatment effect for event e in period t is calculated as:

$$\widehat{\tau}_{et} = \Big(P^{\mathrm{T}}_{et} - \sum_{i \in \mathscr{C}} \hat{\omega}_{ei} P_{it}\Big) - \sum_{s=-T_0}^{-1} \hat{\lambda}_{es}\Big(P^{\mathrm{T}}_{es} - \sum_{i \in \mathscr{L}} \hat{\omega}_{ei} P_{is}\Big)$$

where $\hat{\omega}_{ei}$ are unit weights across control stations and $\hat{\lambda}_{es}$ are time weights across the T_0 pre-treatment months. We calculate this treatment effect for each period both before and after the manager change to verify parallel pre-trends and trace out the evolution of treatment effects over time. We repeat this procedure for every manager-change event, obtaining an event-specific estimate $\widehat{\tau}_{et}$ based solely on its own control station pool and pre-trend.

In the second stage, we regress the estimated effects on traits of the new manager:

$$\widehat{\tau}_{etm} = \alpha + \beta_1 \,\text{CogSkills}_e + \beta_2 \,\text{MentalModels}_e + \mathbf{X}_e + \text{StationControls}_e + \delta_m + \gamma_t + \varepsilon_{etm},$$

where $\mathbf{X}_e$ includes other traits of the new manager (gender, age, experience, noncognitive skills), StationControls includes whether the station is open 24 hours, whether the company rents the station, numbers and types of local competitors, and thirteen location type indicators, δ_m are month fixed effects, and γ_t are turnover event fixed effects. The regression results are displayed in Figure C.5.

To illustrate how the analysis helps rule out confounds related to non-random assignment, we give two examples. (1) Managers with high and low cognitive skills being non-randomly assigned to stations where different local conditions foster different price levels or time trends is not a confound; such effects would be reflected in the pre-period levels and trends of the treated and control stations, and these are differenced out. (2) Managers who are replaced constituting a selected group, e.g., particularly bad managers, is not a confound; if such managers charge systematically different prices over time, this is again captured by the pre-period price levels and trends, which are differenced out.

Figure C.5: Regressions of SDID T.E. on cognitive skills, and mental models

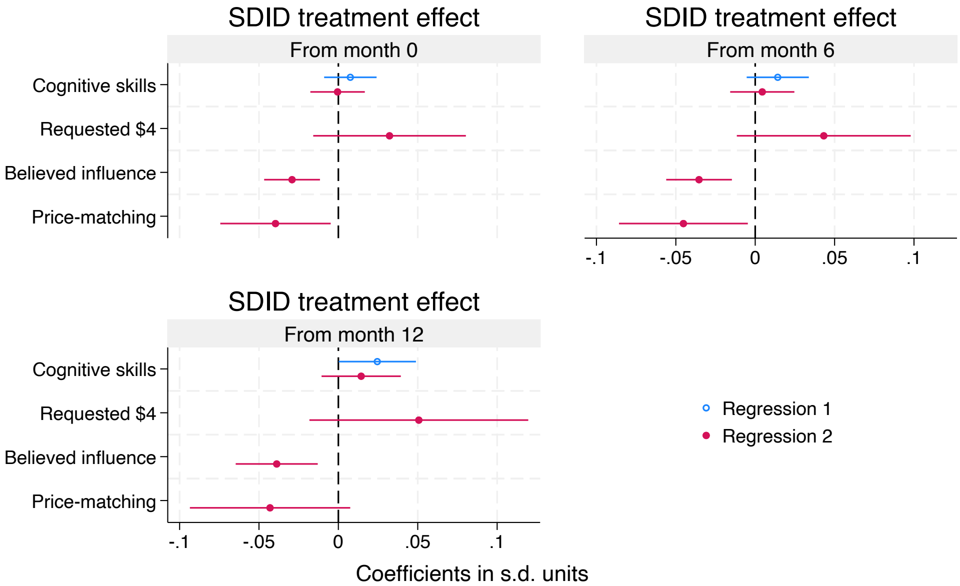

Notes: The figure reports OLS coefficients with 95% confidence intervals. The dependent variable is the SDID treatment effect of the new manager on the price ratio. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, age, gender, experience), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, thirteen location type indicators), and month fixed effects and turnover event fixed effects. Regression 2 includes the three mental models. Robust standard errors are clustered at station level. The different panels consider the entire treatment period, the period starting after 6 months, and the period starting after 12 months. Heteroscedasticity in the estimated dependent variable is addressed by using robust standard errors (Lewis and Linzer, 2005).

C.2.2 Robustness checks for event study

A concern with an estimated dependent variable is that different observations might be measured with more or less noise and influenced by outliers. We check robustness to empirical Bayes shrinkage, which adjusts estimated treatment effects by pulling them toward the population mean, with the degree of shrinkage inversely related to the precision of the estimate to account for sampling error (Kane et al., 2008; Jacob and Lefgren, 2008; Angrist et al., 2017; DellaVigna and Gentzkow, 2019). Intuitively, when a manager's effect is estimated with high noise, we place more weight on the prior that the manager is average, whereas when the effect is precisely estimated, we rely more on the observed data. Our results are robust and, if anything, even stronger after this adjustment (Figure C.6).

Another concern is that the pre-trends for treated and control stations in the SDID analysis could partly be driven by different mechanisms, e.g., if there is some kind of anticipation effect of having a new manager that only affects treated stations. In this case, prices might evolve in different ways after the manager change for treated compared to control stations, but reflecting different types of time trends that continue after the pre-period for one type of station but not the other, rather than a causal effect of the new manager. We address

Figure C.6: Regressions of SDID T.E. on cognitive skills, and mental models after Bayesian shrinkage

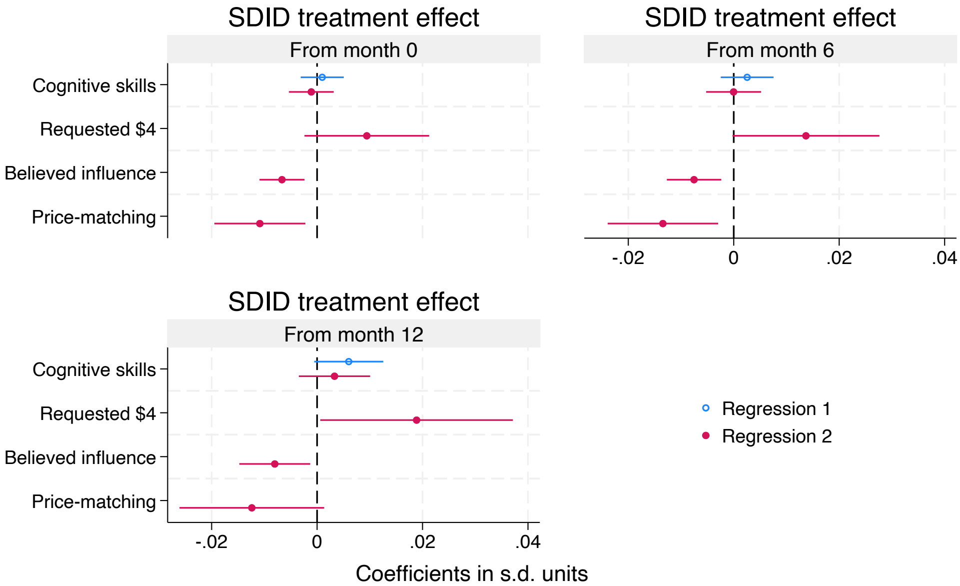

Notes: OLS coefficients with 95% confidence intervals. The dependent variable is the empirical Bayesian shrinkage SDID treatment effect of the new manager on the price ratio. See Figure C.5 for details.

this concern in two ways. One approach is already built into our SDID—using 24 months of pre-treatment data. Trends due to anticipation of a manager move are unlikely to start so far in advance. Since this affects only a small portion of the pre-period, matched control and treatment pre-trends should mainly reflect similar factors. If there is a differential trend for treated stations before the new manager arrives due to anticipated manager change, we would expect to see the pre-period treatment difference deviating from zero in the period leading up to the arrival of the new manager. However, there is little sign of this in our graphs showing treatment differences in the pre-period—the treatment difference is quite stable over the pre-period and notably close to zero right before the manager change (see panels of Figure 7).

As a second way to address this concern, we implemented an alternative difference-in-differences approach. This approach considers only stations that receive new managers and defines a treatment event as receiving a high-skill manager (treatment) versus receiving a low-skill manager (control). This ensures pre-trends in both groups are generated by the same underlying process of manager transition, providing a different type of control group than our main SDID analysis. The approach is also different because it does not seek to estimate a treatment effect for individual stations, which means SDID is not necessary to have a valid control group with parallel pre-trends. We perform a standard difference-in-differences analysis to compare the group of stations that receive above-median ability managers to the group of stations that receive below-median ability managers. While this alternative approach serves as a useful robustness check, we prefer the SDID method for our main analysis for several reasons. First, because of the station-level estimation, SDID allows us to examine the relationship between cognitive skills and pricing in a continuous way across all quintiles, while our robustness check approach can only make binary comparisons between the groups of high-skill and low-skill managers. Second, SDID shows how prices evolve over time for managers with different cognitive skills, allowing us to observe whether high-skill managers raise prices, low-skill managers lower prices, or both. In contrast, the alternative approach only shows the difference between the two groups, making it impossible to determine whether both groups' prices moved in the same direction or opposite directions.

Figure C.7 shows that, using this alternative approach, we find parallel pre-trends when comparing high versus low cognitive skill managers, managers who requested $4 versus those who did not in the money request game, managers with high versus low confidence in their influence over fuel sales, and managers who do versus do not view price-matching as optimal. Turning to the estimated treatment effect, we find that post-change patterns closely resemble our SDID results: high-skill managers and those exhibiting more sophisticated mental models come to charge higher prices over time compared to their counterparts.

Figure C.7: Difference-in-Differences comparing different types of manager transitions

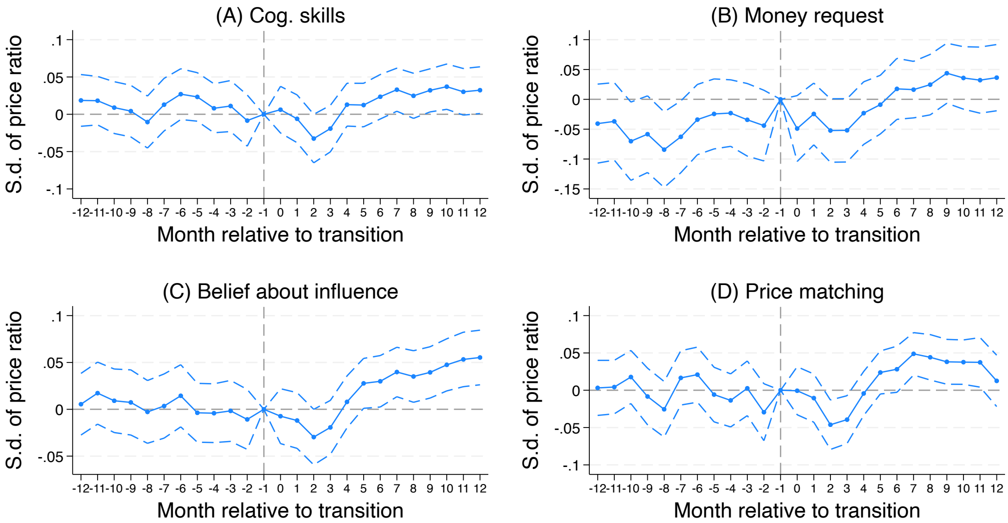

Notes: This figure reports difference-in-differences treatment effects comparing pricing behavior across different types of manager transitions. Unlike SDID, this approach compares transition events to transition events, ensuring pre-trends in both groups reflect the same underlying process of manager change. Each panel uses a different manager characteristic to define treatment and control groups and the more sophisticated managers are defined as the treatment in each case: (A) **Above** median cognitive-skills manager (treatment) vs. **Below** median cognitive-skills manager (control); (B) **Requesting** 4$ (treatment) vs. **Not** in the money request game (control); (C) **Below** median belief about influence over fuel sales (treatment) vs. **Above** median belief about influence over fuel sales (control); (D) **Not** believing in the optimality of the price-matching strategy (treatment) vs. **Believing** in the optimality of the price-matching strategy (control). The y-axis displays standardized price ratios, and the x-axis shows months relative to the manager change. Solid lines represent the difference in mean standardized price ratios between the two groups, with dashed lines showing 95% confidence intervals. A positive treatment effect means that manager in the treatment group charges high fuel prices relative to before than managers in the control group.

C.3 Additional figure and tables

Table C.1: Pricing behavior as a function of cognitive skills and mental models

VARIABLES	Standardized monthly price ratio				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	0.014**	0.014**	0.012*	0.008	0.007
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Requested $4		0.009			0.003
		(0.018)			(0.018)
Believed influence			-0.020***		-0.017***
			(0.006)		(0.006)
Price-matching				-0.074***	-0.071***
				(0.013)	(0.013)
Noncognitive skills	-0.009	-0.009	-0.006	-0.007	-0.004
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Female	0.027**	0.027**	0.027**	0.027**	0.027**
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Age	0.014*	0.014*	0.013	0.011	0.010
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Experience	-0.023***	-0.023***	-0.022***	-0.022***	-0.021***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
# major competitors	-0.057***	-0.057***	-0.056***	-0.055***	-0.055***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
# nearby partner stations	0.037***	0.038***	0.038***	0.038***	0.039***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
# independent competitors	-0.023***	-0.023***	-0.024***	-0.024***	-0.024***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Observations	201,992	201,992	201,992	201,992	201,992
R-squared	0.501	0.501	0.501	0.502	0.502

Notes: *** p<0.01, ** p<0.05, * p<0.1. The table reports coefficients from OLS regression on monthly price ratios. The dependent variable is the standardized station monthly price ratio relative to the price ceiling. In both regressions, we also control station location indicators, whether the station is rented by partner company, open 24 hours, and interacted month and district fixed effects. Robust standard errors are in parentheses, clustered at station level.

Table C.2: The effect of cognitive skills and mental models on SDID treatment effects: Month 0 and onward

VARIABLES	SDID treatment effects				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	0.008	0.006	0.004	0.004	-0.001
	(0.008)	(0.009)	(0.008)	(0.009)	(0.009)
Requested $4		0.039			0.032
		(0.025)			(0.025)
Believed influence			-0.031***		-0.029***
			(0.009)		(0.009)
Price-matching				-0.046***	-0.040**
				(0.018)	(0.018)
Noncognitive skills	-0.014	-0.014	-0.009	-0.013	-0.009
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Female	-0.034**	-0.033*	-0.037**	-0.035**	-0.038**
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Age	0.007	0.008	0.007	0.005	0.006
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Experience	0.008	0.008	0.008	0.008	0.009
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
# of major competitors	-0.018**	-0.018**	-0.017**	-0.016**	-0.015*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
# of nearby partner stations	-0.012	-0.012	-0.012	-0.013	-0.012
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
# of private competitors	0.014**	0.014**	0.013**	0.014**	0.013**
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Observations	81,582	81,582	81,582	81,582	81,582
R-squared	0.019	0.019	0.021	0.020	0.022

Notes: *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the SDID treatment effect of the new manager on the price ratio. The first coefficient in each column is for cognitive skills. Subsequent coefficients are from separate regressions on the given mental model measure, cognitive skills, and controls. Control variables in all regressions include noncognitive skills, experience, age, gender, location indicators, whether the station is rented by partner company, open 24 hours, number of competitors, year-month fixed effects and turnover event fixed effects. Robust standard errors are in parentheses, clustered at station level.

Table C.3: The effect of cognitive skills and mental models on SDID treatment effects: Month 6 and onward

VARIABLES	SDID treatment effects				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	0.014	0.012	0.010	0.010	0.004
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Requested $4		0.051*			0.043
		(0.028)			(0.028)
Believed influence			-0.038***		-0.035***
			(0.011)		(0.011)
Price-matching				-0.054**	-0.045**
				(0.021)	(0.021)
Noncognitive skills	-0.018	-0.018	-0.012	-0.018	-0.012
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Female	-0.043**	-0.042**	-0.047**	-0.045**	-0.047**
	(0.020)	(0.021)	(0.021)	(0.020)	(0.021)
Age	0.005	0.006	0.005	0.003	0.004
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Experience	0.007	0.008	0.007	0.008	0.008
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
# of major competitors	-0.014	-0.014	-0.014	-0.012	-0.012
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
# of nearby partner stations	-0.008	-0.008	-0.008	-0.009	-0.008
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
# of private competitors	0.013*	0.013*	0.012*	0.013*	0.011*
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Observations	48,063	48,063	48,063	48,063	48,063
R-squared	0.020	0.020	0.022	0.021	0.024

Notes: *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the SDID treatment effect of the new manager on the price ratio. The first coefficient in each column is for cognitive skills. Subsequent coefficients are from separate regressions on the given mental model measure, cognitive skills, and controls. Control variables in all regressions include noncognitive skills, experience, age, gender, location indicators, whether the station is rented by partner company, open 24 hours, number of competitors, year-month fixed effects and turnover event fixed effects. Robust standard errors are in parentheses, clustered at station level.

Table C.4: The effect of cognitive skills and mental models on SDID treatment effects: Month 12 and onward

VARIABLES	SDID treatment effects				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	0.025**	0.022*	0.020	0.020	0.014
	(0.012)	(0.012)	(0.012)	(0.013)	(0.013)
Requested $4		0.060*			0.051
		(0.035)			(0.035)
Believed influence			-0.041***		-0.039***
			(0.013)		(0.013)
Price-matching				-0.052**	-0.043*
				(0.026)	(0.026)
Noncognitive skills	-0.025*	-0.025*	-0.019	-0.025*	-0.019
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Female	-0.046*	-0.044*	-0.049*	-0.048*	-0.050*
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Age	0.020	0.020	0.020	0.017	0.018
	(0.017)	(0.017)	(0.016)	(0.017)	(0.017)
Experience	-0.000	0.001	-0.001	0.000	0.001
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
# major competitors	-0.021*	-0.021*	-0.020*	-0.019	-0.019
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
# nearby partner stations	-0.013	-0.012	-0.012	-0.013	-0.012
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
# private competitors	0.018**	0.017**	0.016*	0.017**	0.016*
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Observations	25,844	25,844	25,844	25,844	25,844
R-squared	0.022	0.023	0.025	0.023	0.027

Notes: *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the SDID treatment effect of the new manager on the price ratio. The first coefficient in each column is for cognitive skills. Subsequent coefficients are from separate regressions on the given mental model measure, cognitive skills, and controls. Control variables in all regressions include noncognitive skills, experience, age, gender, location indicators, whether the station is rented by partner company, open 24 hours, number of competitors, year-month fixed effects and turnover event fixed effects. Robust standard errors are in parentheses, clustered at station level.

D Additional results on price wars

D.1 Partner stations' price responses to competitor price cuts

While our main analysis focuses on competitor responses to partner station price cuts, price wars may also begin with competitor-initiated cuts, to which partner managers respond. Cognitive skills may shape these responses, contributing to variation in price war participation. To study this, we define a competitor-initiated price cut as a drop of at least 10 cents following seven days of price stability. We then analyze how partner stations respond, using a difference-in-differences approach that compares their behavior to control stations in the same district whose competitors also maintained stable prices but did not cut prices during the same period.

Figure D.1 shows the results of this analysis. Panel (A) presents the average treatment effects for all managers, while panels (B) and (C) show the effects separately for managers with above-median and below-median cognitive skills. We observe significant heterogeneity in how managers respond to competitor price cuts based on their cognitive skills. Managers with below-median cognitive skills show stronger and more immediate responses to competitor price cuts, reducing their prices by a larger magnitude compared to managers with above-median cognitive skills. This differential pattern is robust to different price cut thresholds, holding for competitor price cuts of 20 cents or more. This finding suggests that bounded rationality not only affects the price level chosen by managers but also influences their strategic responses to competitor actions, potentially amplifying price competition in the market.

Figure D.1: Partner station's price response to competitor price cut

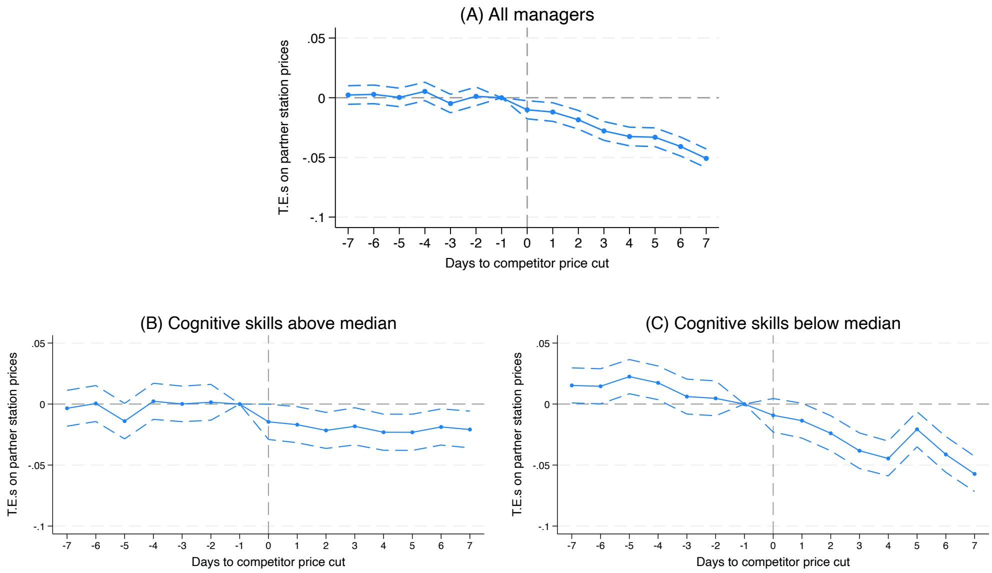

Notes: This figure reports the difference-in-differences treatment effects of partner station price responses to competitor price cuts with 95% confidence intervals. A price cut event is defined as when a competitor station maintains the same price for 7 days and then reduces it by at least 10 cents. The control group consists of stations where the competitor station maintained the same price for 7 days during the same period but did not implement a price cut. The solid line shows point estimates of treatment effects relative to the price cut event (day 0). Panel (A) shows the treatment effects for all managers, while Panel (B) and Panel (C) show the treatment effects for managers with above and below median cognitive skills, respectively.

D.2 Example price war in a local market

Figure D.2: Price competition in a local market

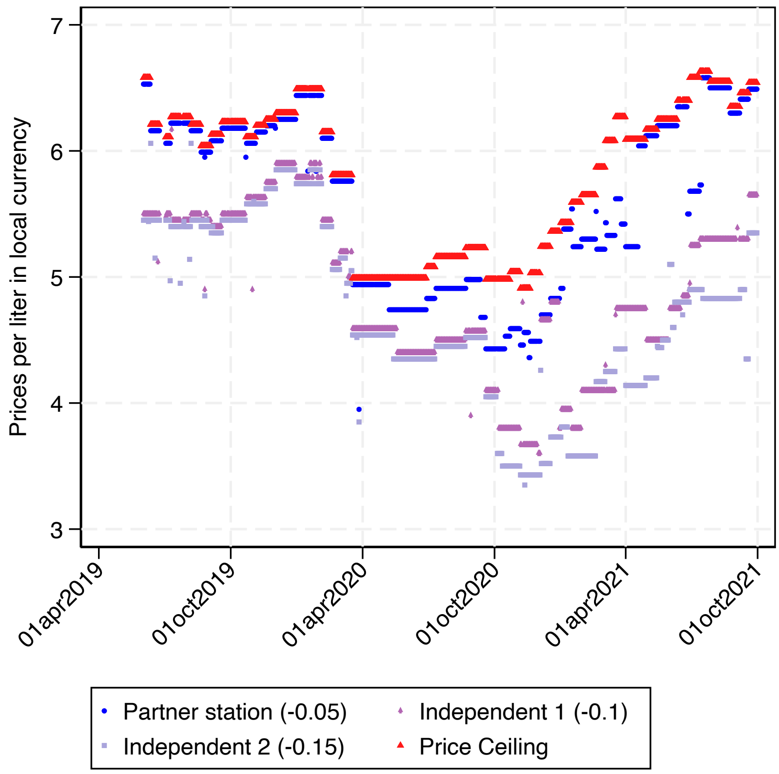

Notes: This figure reports prices per liter in local currency for three stations in a local market from July 1, 2019 to October 1, 2021. The red triangle shows the government-imposed price ceiling. The prices of partner station (-0.05), and two independent competitor stations, Independent 1 (-0.1) and Independent 2 (-0.15), are plotted as dots. The prices for each station are slightly shifted vertically for visual clarity, with the amount of shift indicated in parentheses next to the station's name. The figure shows a price war occurring between October 2020 and February 2021 based on our (conservative) definition of 50 cent price cut or more.

D.3 Regression analysis for price wars

Figure D.3: Number of price wars as a function of manager cognitive skills

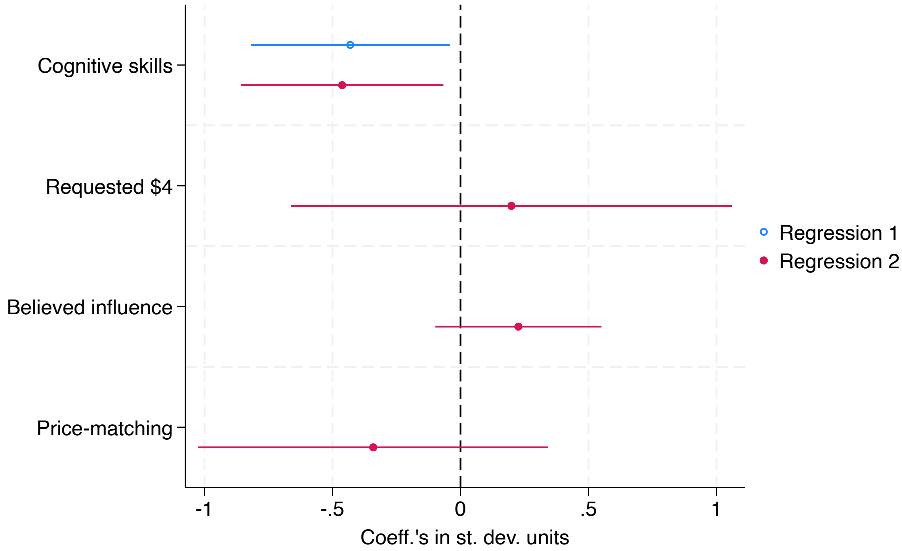

Notes: The figure reports coefficients from negative binomial regressions, with 95% confidence intervals. The dependent variable is the number of price wars. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (noncognitive skills, experience, gender, and age), market structure (numbers and types of local competitors and location type), station characteristics (open 24 hours, and whether the station is rented by partner company), and fuel type and total operation days. Regression 2 also includes the three measures of mental models of competitors.

D.4 Robustness checks for price cut events and price wars

Figure D.4: Competitor's price response to partner station price cut: 20 cents

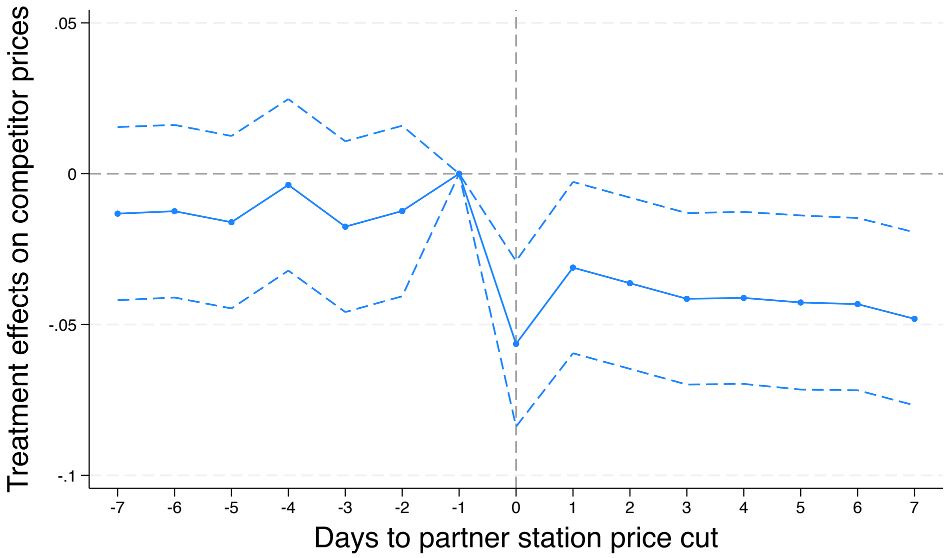

Notes: This figure reports difference-in-differences treatment effects of competitor price responses to partner station price cuts with 95% confidence intervals. Here, A price cut event is defined as when a partner station maintains the same price for 7 days and then reduces it by at least 20 cents. For the rest of the details, see Figure 8.

Figure D.5: Competitor's price response to partner station price cut: 30 cents

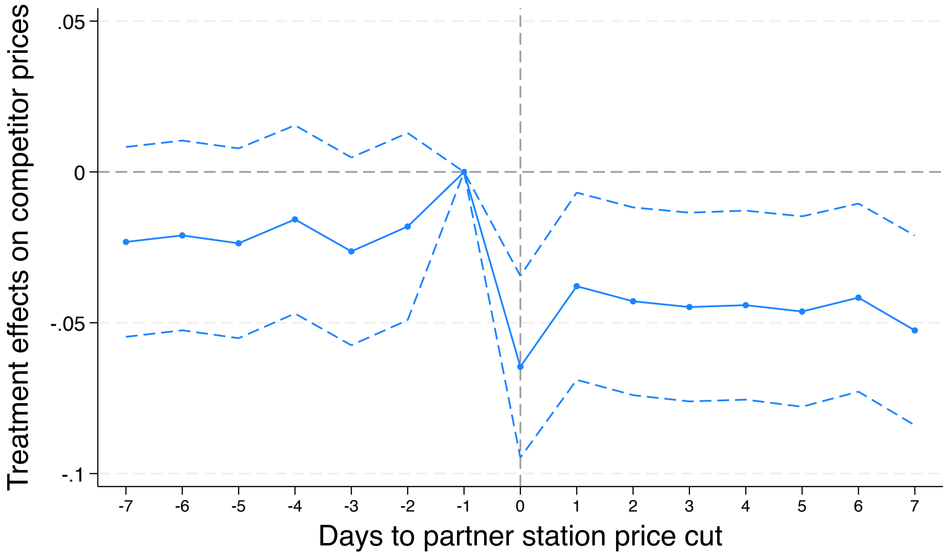

Notes: This figure reports difference-in-differences treatment effects of competitor price responses to partner station price cuts with 95% confidence intervals. Here, A price cut event is defined as when a partner station maintains the same price for 7 days and then reduces it by at least 30 cents. For the rest of the details, see Figure 8.

Figure D.6: Competitor's price response to partner station price cut: SDID

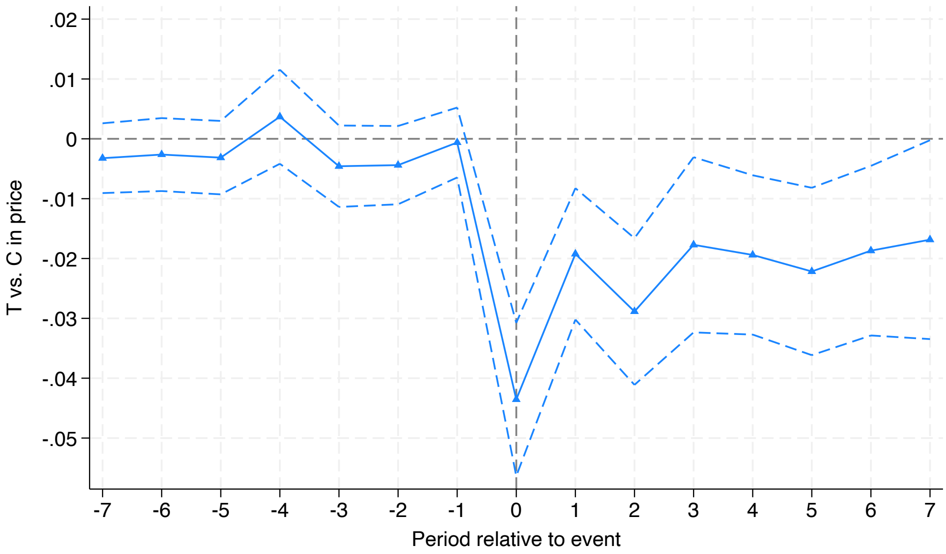

Notes: This figure reports synthetic difference-in-differences treatment effects of competitor price responses to partner station price cuts with 95% confidence intervals. Here the control group for each price cut event is constructed using the Synthetic Difference-in-Differences method. For the rest of the details, see Figure 8.

Figure D.7: Number of price wars as a function of manager cognitive skills in isolated markets

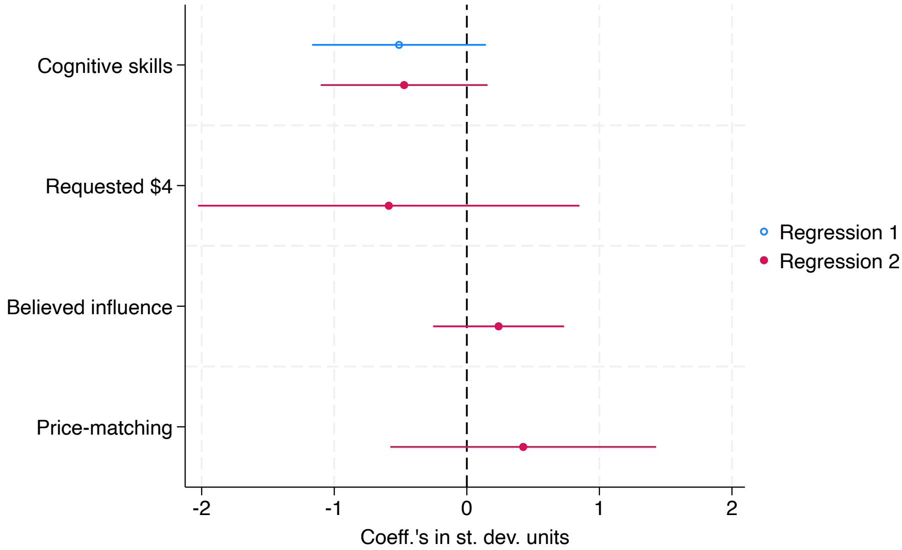

Notes: The sample only includes isolated local markets that do not share common competitors with other local markets. The dependent variable is the number of price wars. For details, see Figure D.3.

Figure D.8: Number of price wars as a function of manager cognitive skills: 25-cents cutoff

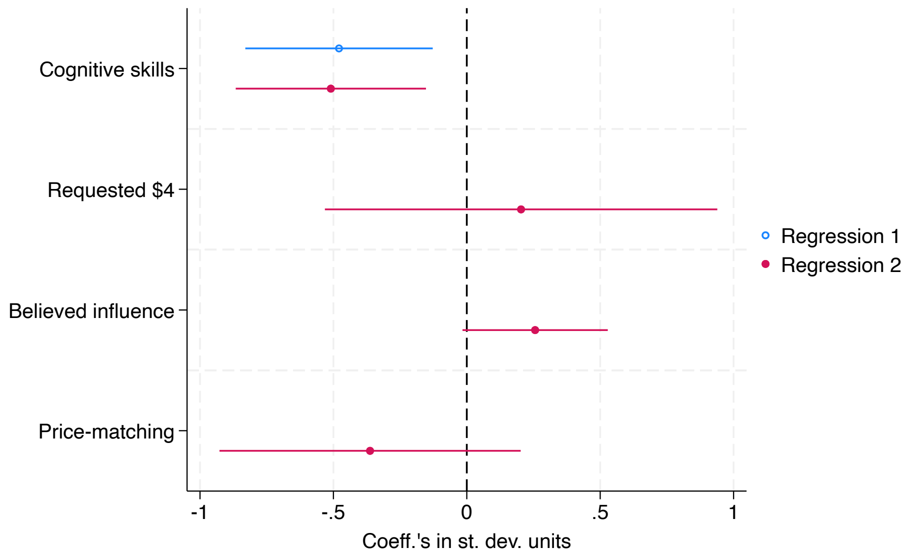

Notes: This table reports the frequency of price wars as a function of manager traits when defining a price wars as "mutual price cuts of at least 25 cents from the price ceiling for a period of 14 days or more." For details, see Figure D.3.

D.5 Additional figures and tables

Table D.1: Number of price wars as a function of manager traits

VARIABLES	Number of price wars				
	(1)	(2)	(3)	(4)	(5)
Cognitive skills	-0.479***	-0.480***	-0.490***	-0.487***	-0.510***
	(0.179)	(0.181)	(0.175)	(0.185)	(0.182)
Requested $4		0.156			0.203
		(0.374)			(0.375)
Believed influence			0.187		0.256*
			(0.145)		(0.139)
Price-matching				-0.208	-0.363
				(0.282)	(0.288)
Noncognitive skills	0.095	0.089	0.081	0.089	0.057
	(0.143)	(0.146)	(0.142)	(0.141)	(0.141)
Female	-0.418	-0.419	-0.474	-0.378	-0.430
	(0.328)	(0.327)	(0.332)	(0.304)	(0.310)
Age	-0.022	-0.026	-0.004	-0.028	-0.013
	(0.224)	(0.227)	(0.211)	(0.228)	(0.219)
Experience	-0.091	-0.091	-0.105	-0.091	-0.111
	(0.188)	(0.189)	(0.188)	(0.186)	(0.188)
# major competitors	0.417**	0.420**	0.439**	0.401**	0.421**
	(0.171)	(0.172)	(0.172)	(0.170)	(0.172)
# nearby partner stations	0.341***	0.347***	0.354***	0.336***	0.357***
	(0.114)	(0.115)	(0.114)	(0.112)	(0.111)
# independent competitors	0.409***	0.405***	0.423***	0.416***	0.440***
	(0.121)	(0.119)	(0.128)	(0.119)	(0.124)
Observations	668	668	668	668	668

Notes: *** p<0.01, ** p<0.05, * p<0.1. This table reports coefficients from negative binomial regression on the number price wars. All specifications controls for noncognitive skills, experience, gender, age, location indicators, whether the station is rented by partner company, station size, number of major competitors, number of nearby own company stations, number of independent competitors, fuel type and total days observed. Robust standard errors are in parentheses.

E Additional results on manager predictions about competitor price responses

Figure E.1: Distributions of managers' predicted competitor price responses by quintile of cog. skills

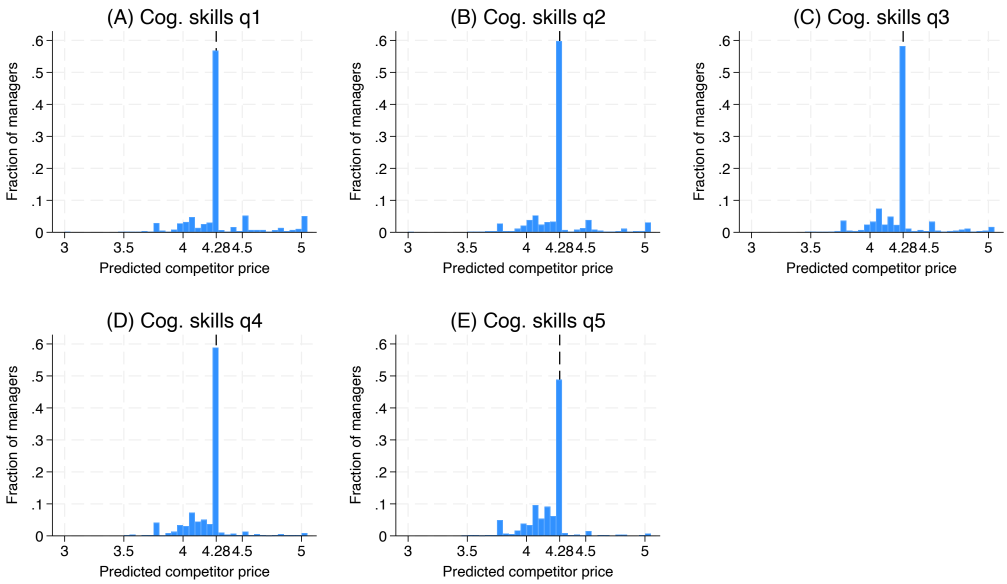

Notes: Distributions of managers' predicted prices of Competitor 1 upon a hypothetical price cut of a partner station by quintiles of manager cognitive skills (quintile 5 is the highest). Here, 4.28 is the actual price of Competitor 1 without the hypothetical price cut.

Figure E.2: Cognitive skills and predicted competitor 2 price responses

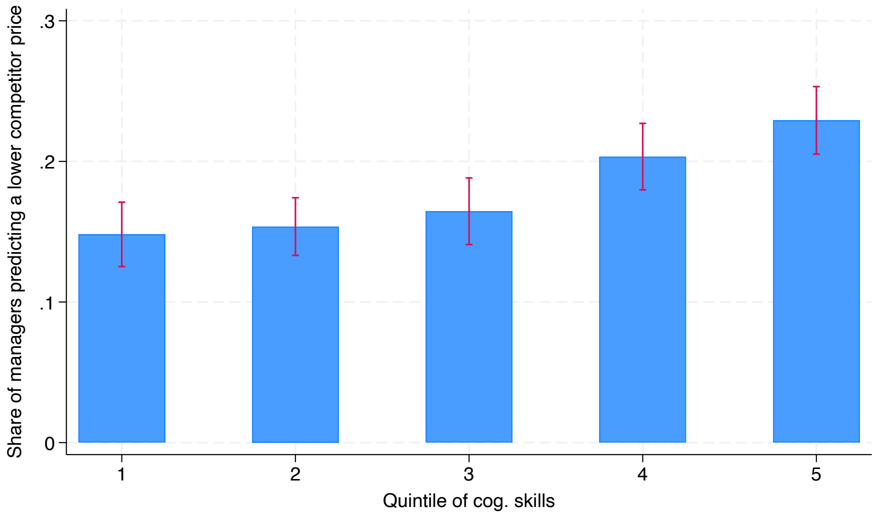

Notes: This figure reports the share of managers predicting a lower Competitor 2 price in response to a hypothetical price cut of a partner station by quintiles of cognitive skills (quintile 5 is the highest).

Figure E.3: Cognitive skills, money request game, and predicted lower competitor prices

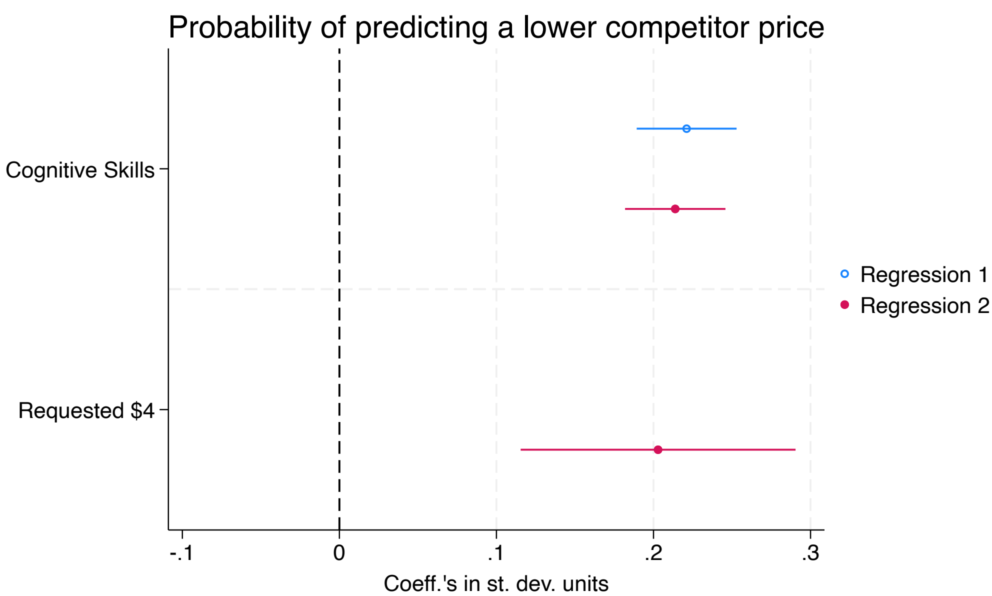

Notes: This figure reports marginal effects from Probit regressions, with 95% confidence intervals. The dependent variable equals 1 if the manager predicted a lower Competitor 1 price following the hypothetical price cut by our partner station. Regression 1 reports the coefficient for cognitive skills but also controls for other manager traits (experience, gender, and age) and station characteristics (open 24 hours, numbers and types of local competitors, district fixed effects). Regression 2 adds an indicator for whether the manager requested $4 in the money request game. The sample consists of managers participating in the third survey wave, with dependent and independent variables measured in the third wave.

Figure E.4: Predicted lower competitor prices and mental models

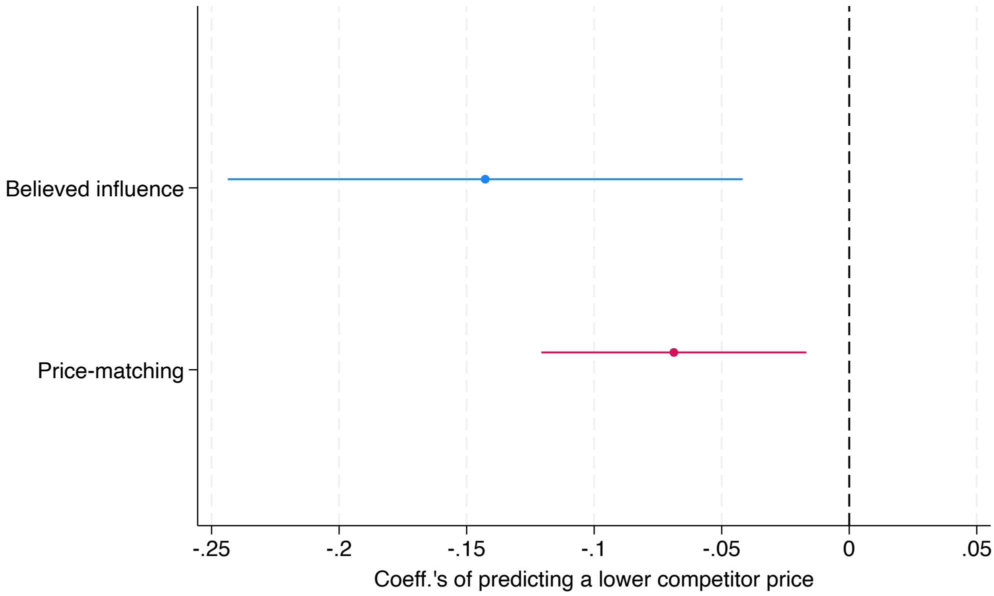

Notes: This figure reports OLS regression coefficients, with 95% confidence intervals. The independent variable equals 1 if the manager predicts a lower Competitor 1 price following the hypothetical price cut by our partner station. The first coefficient shows its effect on believing managers can influence fuel sales. The second coefficient shows its effect on believing price-matching is optimal. The sample consists of managers participating in both the first and third survey waves, with mental model measures from the first wave and the prediction task from the third wave. Both regressions controls for manager traits (cognitive skills, noncognitive skills, experience, gender, and age) and station characteristics (open 24 hours, numbers and types of local competitors, district fixed effects).

F Additional results on profits and welfare

F.1 Estimating demand parameters for welfare analysis

For our welfare calculations, we focus on the average station rather than conducting station-by-station analysis. We adopt this approach due to two data limitations. First, our estimates of how cognitive skills affect pricing come from regressions that identify the average effect across all stations. Estimating station-specific effects would require additional identifying assumptions about the heterogeneous impact of cognitive skills across different stations. Second, our measure of competitor responsiveness θ is estimated from price-cut events aggregated across all stations. Obtaining station-specific estimates of θ would require observing frequent price-cutting events at each individual station, which is not feasible given the relatively infrequent nature of such events.

Our demand function for the average station implies that the quantity sold evolves according to

$$\frac{d\bar{q}}{d\bar{p}} = \bar{b} + \bar{c}\,\bar{\theta},$$

where $\bar{b}$ captures the *direct* response of demand to the station's own price, $\bar{c}$ measures the *cross-price* response to competitors' prices, and $\bar{\theta} \in [0,1]$ is the share of the station's own price change that competitors are expected to match on average.

Identification strategy using dual cost shifters We exploit two institutional features for identification. First, the government price ceiling is updated every ten days following a predetermined formula indexed to world oil prices. Second, our partner company's internal accounting rule mandates that marginal cost moves with the current ceiling *one-for-one*, which also means the previous price ceiling has no effect on current price controlling for the current ceiling. Independent competitors purchase their refined oil from independent refineries and thus their marginal costs adjust to the price ceiling with a delay. This means that the marginal cost of competitors is affected by both the current price ceiling and the price ceiling in the previous cycle.

Let $Ceiling_t$ denote the ceiling that is in force during the current pricing cycle t, and $Ceiling_{t-1}$ the ceiling from the previous cycle $t-1$. These two cost shifters affect our partner stations and competitors differentially, providing the variation needed for identification. We treat $Ceiling_t$ and $Ceiling_{t-1}$ as a pair of excluded instruments for the two endogenous price regressors: the partner station's own price and the market-average competitor price. The estimating equation for station s and fuel product j is

$$q_{sjt} = a_{sj} + b\,p_{sjt} + c\,p_{-sjt} + \delta_{sjm} + \phi_{sjw} + \psi_{sjy} + \eta_{sjt},$$

where q_{sjt} is the quantity sold in period t at station s of product j, p_{sjt} and p_{-sjt} are the corresponding own and competitor prices respectively, a_{sj} represents the station-product fixed effect, δ_{sjm} represents month fixed effects, ϕ_{sjw} represents day-of-week fixed effects, ψ_{sjy} represents year fixed effects, and η_{sjt} is an error term.

The first-stage regressions confirm the expected differential impact of our cost shifters. For partner stations' own prices, the coefficient on $Ceiling_t$ is 0.98 (s.e. 0.001), while the coefficient on $Ceiling_{t-1}$ is only 0.01 (s.e. 0.001), confirming the near one-for-one pass-through of current costs and negligible influence of lagged ceilings. For competitor prices, the pattern is different: the coefficient on $Ceiling_t$ is 0.73 (s.e. 0.001), indicating substantial but incomplete pass-through, while the coefficient on $Ceiling_{t-1}$ is 0.20 (s.e. 0.001), reflecting the persistent influence of previous cycle costs. This differential response pattern validates our identification strategy.

The second-stage coefficients from the instrumental variables regression yield the following elasticities:

$$\bar{b}\,\frac{\bar{p}}{\bar{q}} \;=\; -2.07, \qquad \bar{c}\,\frac{\bar{p}}{\bar{q}} \;=\; 1.15,$$

Our event study of competitor responses to price cuts that are not driven by a change in price ceiling (Section 5) provides an independent estimate of the average competitor response at $\bar{\theta} = 0.38$. This strategic response parameter, combined with our demand parameter estimates, yields the effective price elasticity when a manager decides to change their price keeping the price ceiling fixed:

$$\frac{d\bar{q}}{d\bar{p}}\,\frac{\bar{p}}{\bar{q}} = (\bar{b} + \bar{c}\,\bar{\theta})\frac{\bar{p}}{\bar{q}} = -2.07 + 0.38 \times 1.15 = -1.63.$$

F.1.1 Suggestive evidence for heterogeneity in station-level demand elasticities

To show that elasticity varies substantially across stations, implying that optimal prices should differ across locations, we estimate station-level demand parameters. A key limitation is that we cannot observe station-level competitor responsiveness θ_{sj} due to the infrequency of price-cutting events at individual stations. Instead, we estimate station-level elasticity when the price ceiling changes, which captures how competitors respond to cost shocks rather than strategic price cuts.

Because we do not need to disentangle direct response of demand and cross-price response for this analysis, we adopt a *single-instrument* strategy that omits competitor price in the estimation and assume instead that it responds to the price ceiling with a parameter δ_{sj}:

$$p_{-sjt} = \beta_{sj} + \delta_{sj}\,Ceiling_t + \varepsilon_{-sjt},$$

where $\delta_{sj} \in [0, 1]$ is the *local pass-through rate* linking competitors' prices to the current ceiling. This parameter is allowed to differ by station s and product j. Substituting this expression into the baseline demand equation yields a *residual-demand* slope

$$\frac{dq_{sjt}}{dp_{sjt}} = b_{sj} + \delta_{sj}\, c_{sj},$$

which we estimate using $Ceiling_t$ as the sole excluded instrument for p_{sjt}.

For each station-product pair we estimate

$$q_{sjt} = a_{sj} + \beta_{sj}\, p_{sjt} + \delta_{sjm} + \phi_{sjw} + \psi_{sjy} + \eta_{sjt},$$

where p_{sjt} is instrumented with $Ceiling_t$ and the fixed effects structure mirrors the market-level regression. The 2SLS slope $\hat{\beta}_{sj}$ estimates $b_{sj} + \delta_{sj} c_{sj}$, the elasticity that a manager faces when the price ceiling changes and competitors respond with pass-through rate δ_{sj}.

This measure likely provides a lower bound on the true heterogeneity in elasticities across stations. Cost pass-through rates δ_{sj} are probably similar among independent competitors, as they all face similar wholesale cost structures. In contrast, strategic responsiveness θ_{sj} to competitor price cuts likely varies more across local markets, depending on factors such as the competitive intensity and manager sophistication at competing stations.

Figure F.1: Estimated demand elasticities

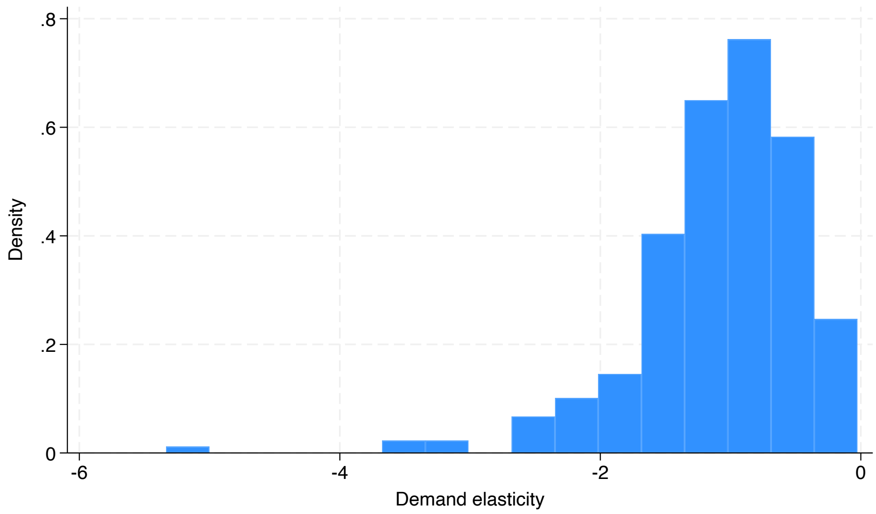

Notes: Distributions of estimated station-level price elasticities. These elasticities capture the combined effect $b_{sj} + \delta_{sj} c_{sj}$. The elasticities are estimated by regressing daily sales volume on price, instrumenting price with the price ceiling.

Nevertheless, even this conservative measure demonstrates substantial variation in demand conditions across stations. As shown in Figure F.1, elasticity estimates range from below -5 to above -1, with most stations concentrated between -3 and -1. This heterogeneity suggests that determining optimal prices for each station might be challenging.

F.2 Details on calculations of producer surplus, consumer surplus, deadweight loss, and markups

Using the linear demand function estimated in our model, we calculate the welfare implications of differences in pricing strategies between high-skill and low-skill managers for an average station in the district with daily data. Below, we detail our approach for computing each welfare measure.

Let p_h and p_l denote the prices charged by high-skill and low-skill managers, respectively, with $p_h > p_l$ based on our findings. Let q_h and q_l represent the corresponding quantities sold at these prices. Recall that the demand function is given by:

$$q_i = a + b_i p_i + c_i p_{-i} + g(z).$$

The derivative of quantity with respect to own price equals:

$$\frac{dq_i}{dp_i} = b_i + c_i \theta_i$$

Figure F.2: Welfare comparison: high-skill vs. low-skill managers

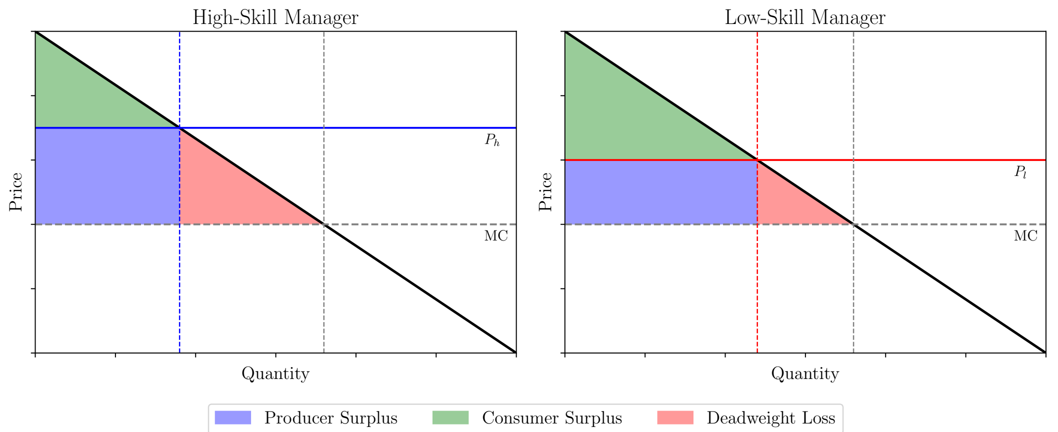

Notes: This figure reports the welfare comparison between high-skill and low-skill gas station managers using simulated data. The left panel shows welfare components for a high-skill manager who charges a higher price (P_h), while the right panel shows welfare components for a low-skill manager who charges a lower price (P_l). This visualization uses hypothetical demand curves, prices, and marginal costs for illustrative purposes only and does not correspond to our actual empirical estimates.

Producer Surplus (PS). Producer surplus is calculated as revenue minus variable costs:

$$PS = p \cdot q - MC \cdot q = (p - MC) \cdot q$$

The difference in producer surplus between high-skill and low-skill managers is:

$$\Delta PS = PS_h - PS_l = (p_h - MC) \cdot q_h - (p_l - MC) \cdot q_l$$

Here we set p_h equal to the average price $\bar{p}$, and calculate p_l as p_h minus twice the standardized effect of cognitive skills on fuel prices, capturing the price effect of replacing one of the highest skilled managers with one of the lowest. Given the demand function and the calculated slope, we then compute the corresponding sales under the two prices. The percentage change in producer surplus is calculated as $\frac{\Delta PS}{PS_h}$. The difference in producer surplus depends on the marginal cost, and we calculate it for the whole range of plausible marginal costs.

Consumer Surplus (CS). With linear demand, consumer surplus is the area of the triangle under the demand curve and above the price:

$$CS = \frac{1}{2} \cdot (p^{max} - p) \cdot q$$

where p^{max} is the price at which quantity demanded becomes zero. The difference in consumer surplus is:

$$\Delta CS = CS_l - CS_h = \frac{1}{2} \cdot [(p^{max} - p_l) \cdot q_l - (p^{max} - p_h) \cdot q_h]$$

For our calculations, we can simplify this using the linear demand relationship:

$$\Delta CS = \frac{1}{2} \cdot (p_h - p_l) \cdot (q_h + q_l)$$

The percentage change in consumer surplus is calculated as $\frac{\Delta CS}{CS_h}$. The difference in consumer surplus does not depend on the marginal cost, and the lower prices charged by the low-skill managers increases the conumser surplus, as shown in Figure F.2.

Deadweight Loss (DWL). Deadweight loss measures the welfare loss from pricing above marginal cost:

$$DWL = \frac{1}{2} \cdot (p - MC) \cdot (q^{MC} - q)$$

where q^{MC} is the quantity that would be sold if price equaled marginal cost. The difference in deadweight loss is:

$$\Delta DWL = DWL_h - DWL_l = \frac{1}{2} \cdot [(p_h - MC) \cdot (q^{MC} - q_h) - (p_l - MC) \cdot (q^{MC} - q_l)]$$

The percentage change in DWL is calculated as $\frac{\Delta DWL}{DWL_h}$. The magnitude of difference in DWL depends on the marginal cost, and we calculate the change in DWL across the whole range of plausible marginal cost. However, there is a reduction in DWL going from the high-skill manager to the low-skill manager regardless of the marginal cost, as shown in Figure F.2.

Price Markup. The price markup is calculated as:

$$Markup = \frac{p - MC}{p}$$

The difference in markup is:

$$\Delta Markup = Markup_h - Markup_l = \frac{p_h - MC}{p_h} - \frac{p_l - MC}{p_l}$$

The percentage change in markup is calculated as $\frac{\Delta Markup}{Markup_h}$. Similar to DWL, the magnitude of difference in markup depends on the marginal cost, but the markup is always higher with high-skill managers than with low-skill managers.

F.3 Additional figures and tables

Table F.1: Profits as a function of manager and station characteristics and instrumented price

VARIABLES	Total profit			Fuel profit			Nonfuel profit		
	OLS	OLS	IV	OLS	OLS	IV	OLS	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Price ratio		0.096***	0.200**		0.106***	0.225**		-0.021***	-0.068
		(0.005)	(0.086)		(0.005)	(0.090)		(0.005)	(0.086)
Cognitive skills	0.014***	0.012***	0.011**	0.012***	0.011**	0.009**	0.018***	0.019***	0.019***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)
Noncognitive skills	0.004	0.005	0.006*	0.004	0.005	0.006*	0.003	0.003	0.003
	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Female	0.016*	0.013	0.011	0.014*	0.012	0.009	0.022**	0.023**	0.024**
	(0.008)	(0.008)	(0.008)	(0.009)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)
Age	-0.001	-0.003	-0.004	-0.002	-0.003	-0.005	-0.000	0.000	0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Experience	0.022***	0.024***	0.027***	0.021***	0.023***	0.026***	0.022***	0.022***	0.021***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)
Market share	0.031***	0.029***	0.028***	0.031***	0.029***	0.027***	0.022***	0.022***	0.023***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)
# major competitors	-0.007*	-0.001	0.004	-0.008**	-0.002	0.004	0.006	0.005	0.003
	(0.004)	(0.004)	(0.006)	(0.004)	(0.004)	(0.006)	(0.005)	(0.005)	(0.007)
# nearby partner stations	0.012***	0.008**	0.004	0.012***	0.008*	0.003	0.009**	0.010**	0.012**
	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)
# independent competitors	0.000	0.002	0.005	-0.000	0.002	0.005	0.004	0.004	0.003
	(0.002)	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
Observations	197,156	197,156	197,156	197,156	197,156	197,156	197,156	197,156	197,156
R-squared	0.031	0.037	0.030	0.029	0.035	0.027	0.018	0.018	0.017

Notes: *** p<0.01, ** p<0.05, * p<0.1. The table report coefficients from 2SLS regression, instrumenting price with mental models. Control variables include cognitive skills, noncognitive skills, experience, age, gender, location indicators, whether the station is rented by partner company, open 24 hours, number of major competitor stations, number of own company nearby stations, number of independent competitor stations, market share, and interacted month and district fixed effects. Results are from monthly data. The first stage F-value is 443.27. Robust standard errors are in parentheses, clustered on station and month level.

H Narrative rubric

1. Background Information

This guide aims to help understand how to classify responses from gas station managers to the question, “Some gas station managers consistently have higher fuel profits; how do they achieve this?”

Since profit = profit margin × sales volume, we divide the managers’ responses into two main categories: Profit Margin and Sales Volume. Responses that cannot be classified into these two categories are placed in “Others.”

2. Classification Overview

- Main Categories: Profit Margin, Sales Volume, Others.
- Profit Margin Subcategories: Higher Prices, High-Profit Fuel Ratio, Targeted Discounts/Differential Pricing, Cost Reduction, Market Conditions, Missing
- Sales Volume Subcategories: Customer Development, Location/Market, Low Price/Discounts, Missing
- Other Subcategories: Manager Capability, Effort, Don’t Know/Unclassifiable

3. Classification definition and examples

Category	Subcategory	Definition and key terms	Examples
Profit Margin	Higher Prices	The "high-price" subcategory focuses on the strategy of gas stations increasing the profitability of oil products by raising fuel prices, reducing discounts and promotional activities, etc. Keywords: reduce marketing expenses, eliminate discounts, price positioning rate, decrease discounts, raise oil product prices, reduce promotional activities, maintain high prices, avoid price wars, stabilize customers, increase oil product gross profit margin.	• "Reduce marketing expenses" • "Cancel discount mechanism" • "The station focuses on oil product profitability and rarely reduces prices" • "Increase price realization rate" • "Utilize personal skills to market to customers and minimize their use of company's promotional policies"
	High-profit oil product ratio	Indicates increasing the proportion of high-profit oil products in the sales mix. Keywords: High gross profit, premium-grade products	• "Sell more high-margin products" • "Sell more high-margin oil products"

	Targeted discounts or differential pricing	The "Targeted Marketing" subcategory emphasizes implementing personalized promotional activities and pricing strategies for different customer groups or individual customers through precise marketing strategies, without universal applicability. Keywords: One-to-one pricing, One-to-one strategy, Precision marketing, Point-to-point promotion, Customer segmentation management, Differentiated marketing, Buy-more-get-more, Customer segmentation, Tiered activities, Price grading, Market research, Precision promotion.	• "Implementing different promotions for different customers" • "One price per customer" • "Point-to-point promotion" • "Precision marketing approach" • "For customer segmentation management, promotional offers vary" • "Refine customers, price grading" • "Differential marketing activities, carry out promotions for those activities, and prioritize profitability"

	Lower Costs	The “Lower Costs” subcategory emphasizes increasing the profitability of gas stations by reducing operational costs, optimizing resource utilization, and controlling daily expenses. This includes but is not limited to reducing operating expenses, energy conservation, controlling labor costs, and optimizing on-site expenditure. Keywords: Cost-saving, expenditure reduction, cost control, optimizing workforce, energy conservation, cost management, electricity savings, labor cost control, daily consumption control, cost reduction, expenditure reduction, saving and spending optimization, minimizing gas station expenses. Note: “marketing expense” is not a cost, it means lowering prices.	• “Save on-site expenses to increase oil product gross profit” • “Reduce operating expenses” • “Source optimization and flow control” • “Optimize labor utilization, reduce consumption, and save energy” • “Maximize gross profit, minimize expenses” • “Save electricity” • “Energy conservation and consumption reduction”

	Market Conditions	Under favorable market conditions (such as advantageous geographical location, low competition, etc.), gas stations are able to maintain or increase oil prices, thereby directly increasing profit margins. Keywords: High market demand, no need for price reduction, high gross profit demand, market advantage, low competition	• “High demand in surrounding markets for high-margin oil products” • “In a good market, gas stations do not need to lower prices.” • “In a good market, good profits are realized.”
	Missing	Only mention increasing profit rate, without mentioning how to increase profit rate. Keywords: Profit rate, gross profit rate. Note: If only profit is mentioned, it should be classified as ’unknown’ in ’Others’. Here, pay attention to the distinction between profit and profit rate.	• “Increase oil product gross profit margin”

Sales Volume	Customer development	Developing customers refers to various proactive strategies and actions taken by gas station managers to attract new customers, maintain existing customer relationships, or restore past customers, thereby increasing sales volume and enhancing customer loyalty. Keywords: Developing new customers, going out to run business, customer maintenance, actively expanding the market, increasing signed customers, customer relationship management, personalized services, market research, promotional activities, online marketing, face-to-face communication, enhancing brand value, establishing customer trust, optimizing service quality, increasing customer loyalty, market expansion, vigorously grasping customers. Note: Verb + customer denotes customer development, while customer + adjective denotes location or market	• "Developed new customers" • "Go out to run business" • "Maintain customers well" • "Actively explore the market" • "Increase the number of signed customers"

	Location or Market	Market conditions or location factors have a positive impact on gas station sales, such as ample customer base and a high volume of surrounding customers, by attracting more customers to increase sales. Keywords: Abundant customers, sufficient customer sources, lack of competition, excellent customer resources, favorable surrounding market, market determination, market demand, advantageous market, market opportunities, geographic location, entry rate, city entrances and exits, highway entrances, corporate units, advantages of gas stations, surrounding markets, local conditions, high-quality customers, favorable market, high market demand. Note: Verb + customer denotes customer development, while customer + adjective denotes location/market	• "Abundant customers" • "Sufficient customer sources" • "Lack of competition" • "Excellent customer resources" • "Favorable surrounding market"

	Low Prices or Discounts	The "Low Price" subcategory focuses on increasing the sales volume of gas stations by providing promotional activities, lowering fuel prices, or implementing a strategy of low-profit, high-sales. These practices aim to compensate for the decrease in profit per transaction by increasing quantity, thereby overall enhancing or maintaining profit levels. Keywords: Low-profit, high-sales, promotional activities, price discounts, promotions, price reductions, fuel promotions, discounts, direct reductions, hosting events, promotional intensity, continuous promotions, member day discounts.	• "Gas stations hold promotional activities" • "Long-term promotion of oil products" • "Discounts per ton of oil" • "Discounted price reduction" • "Oil products have promotions" • "There are quite a few promotional activities"
	Missing	Only mention increasing sales volume without specifying how to increase sales volume. Keywords: Sales volume, Increment	• "Increment" • "Increase Sales volume"

Others	Managers' capability	The role of a gas station manager's personal abilities, management skills, business insight, execution, and leadership in increasing the profit rate and sales volume of the gas station. This includes but is not limited to operational management, personal charm, family support, understanding of business strategies, and luck. Keywords: Management ability, personal ability, family strength, leadership appreciation, business insight, management skills, execution, entrepreneurial mindset, astute and capable, observant, excellence, ability surpasses all, affinity.	• "Strong management ability" • "Personal ability and family strength are stronger than others." • "Skilled in management" • "Strong capability, agile mind" • "Proper management, proficient in business" • "Strong execution, excellent"

	Effort	The "Effort" subcategory emphasizes the important role of the station manager and staff's work attitude, sense of responsibility, and level of commitment in improving the profit rate and sales volume of the gas station. This includes enthusiasm for work, attentiveness to customer service, and wholehearted dedication to operational management. Keywords: Effort and reward, diligence, conscientiousness, seriousness, wholehearted dedication, down-to-earth, facing challenges head-on, diligence, striving, diligence, sense of responsibility, initiative, leading by example, active effort, treating the station as home, reaping what you sow.	• "Others' efforts are equal to returns." • "Done with care." • "Serious and responsible." • "Wholeheartedly devoted to the gas station." • "Conscientious and proactive."

	Understanding the Market	The "Understanding the Market" subcategory encompasses a station manager's understanding of the market environment, keen insight into market dynamics, and the ability to formulate effective marketing strategies and response mechanisms based on this information. Keywords: Understand the market, grasp the market, market knowledge, market understanding	• "Keep a close eye on market environment, understand market dynamics, establish response mechanisms, flexibly use promotional policies" • "Understand market demand, master customer dynamics" • "Understanding the market and surrounding environment allows clear analysis of customer needs and development of reasonable policies for different customer groups" • "Good grasp of oil product market, able to adjust marketing strategies in a timely manner to attract different customer groups"

	Analytical Ability	The "Analytical Ability" subcategory reflects the financial analytical capabilities that station managers possess when operating gas stations, including budgeting skills, cost and profit calculations, and the ability to manage accounts precisely. Keywords: Good at analysis, financial analysis, budgeting, accounting, calculation	• "Station manager is good at analysis, maximizing profits" • "Has business acumen, knows how to keep accounts, manages station accounts well" • "Focuses on data, calculates carefully" • "Calculates to win in advance, clearly distinguishes customers"
	Don't know or unrecognizable	Answers that cannot be clearly categorized into any of the above subcategories and do not make sense. Keywords: Achieved, endless, no opinion, not very familiar, none, thank you, don't understand, redundant, have a clue, not something I consider, gold shines everywhere it goes, unclear, no opinion, very good, pretty good, impressive, useless, don't understand, this question is quite complicated.	• "No opinion" • "Not very familiar" • "Profit is the lifeblood of a company" • "Don't understand" • "Redundant"

4. Operational Steps

- Determine Subcategories: Attempt to classify responses into specific subcategories, then code as 0 or 1. Note the difference between "Missing (Profit Margin)" and "Missing (Sales Volume)."

- Handling Multiple Categories: If a response involves multiple categories (e.g., "Using personal skills to market to customers, minimizing their use of company promotions"), and the relationship is parallel, then it should be marked under both relevant categories. However, if the factors are causally related, such as "The gas station's market is strong, so no need to lower prices," only mark "Market Conditions (Profit Margin)" and not "High Price."

I Station manager survey first wave

[The following questions are about the gas station:]

Q1. The name of the gas station (Please specify the full name of the station).

____

Q2. The ID of your station.

____

Q3. Please indicate, how many employees are employed in the management team at your station: (number)

____

Q4. Please indicate how many employees at your gas station are hired through a third-party company: (number)

____

Q5. How many hours does your gas station operate on a weekday?

hours__________

Q6. What is the average monthly wage for an average employee in your station? Please enter the monthly amount in the local currency:

____

Q7. Compared to other nearby gas stations, what is the salary level of the average employee at this site? Please answer with a number from 1 to 5, where 3 means the salary level of employees at this site is equal to the average level of nearby gas stations, 5 represents the salary of the staff at this station is higher than most nearby stations, 1 represents the salary of the staff at this station is lower than most nearby stations, the larger the number, the higher the income of the staff at this station (compared to nearby stations) [1-5]

____

Q8. How long does it take for you to hire a new employee if you would like to hire one today? Please enter the number of weeks:

__________ weeks

Q9. What is the market share of your gas station in oil sales? Please answer with a number between 0 and 100.

For example, 75 means 75% market share.

____

[The following questions are about working experience at the gas station]

Q10. What is your current role

1. Station Manager
2. Gas Station Employee
3. Others__________, please specify__________

[Q10≠1] Q11. The following question can only be answered by the station manager. Please ask your station manager to answer the questions from now on.
[Q10≠1] Q12. Is it the station manager who is answering the survey?

1. Yes
2. No

[Q10≠1] [Q12=2] Q13. If the station manager is not able to answer the survey, please make sure to consult your station manager on how to answer each question. When we refer to "you", we mean "the station manager".
Q14. Since when are you working at this company? Please enter the starting date:
Month __________ Year__________
Q15. Since when are you working at this gas station? Please enter the starting date:
Month __________ Year__________
Q16. How many years have you worked as a station manager (at this or other stations of the company)? Please enter the number of the years
__________ Years
Q17. Did you work at another station at this company before the current one?

1. Yes
2. No

[Q17=1] Q18. What is your role in that station? station manager, station employee, others: please specify

1. Station Manager
2. Station Employee
3. Others:__________

[Q17=1] Q19. What is the ID of that station? __________ What is the name of that station?

____
Q20. How many years have you worked in this previous station? [Please enter the number of the years]
__________ Years
Q21. How many years have you worked as a station manager (at this or other stations of the company)? [Please enter the number of the years]
__________ Years
[The following questions are about demographics]:

Q22. How old are you? [Please enter your age]

Q23. What is your gender?

1. male
2. female

Q24. What is your highest level of education?

1. High School or Below
2. Junior College
3. Undergraduate
4. Graduate or above

[Q24>1] Q25. What is your major at the university?

1. Philosophy
2. Economics
3. Management
4. Law
5. Education
6. Literature
7. History
8. Science
9. Engineering
10. Agriculture
11. Medicine
12. Military
13. Computer
14. others__________

Q26. Where do you live on weekdays (Monday to Friday)?

1. Home
2. Dormitory provided by the company
3. Others__________

[The following questions are about management practices:]

Q27. Compared to external factors such as the location of the gas station and the amount of competitors, in your opinion, to what extent can managers influence the performance of a station? [0-100] Please choose a number on a scale between 0 and 100, where 0 means managers have no influence on how well a station operates, 100 means how well a station performs is solely determined by its manager.

____

Q28. Compared to objective factors such as the location of the gas station and the amount of competitors, in your opinion, to what extent can managers influence the sales of oil products of a station? [0-100] Please choose a number on a scale between 0 and 100, where 0 means managers have no influence on **oil products sales**, 100 means the amount of **oil products sales** is solely determined by the manager.

____

Q29. Compared to objective factors such as the location of the gas station and the amount of competitors, in your opinion, to what extent can managers influence the non-oil profits of a station? [0-100] Please choose a number on a scale between 0 and 100, where 0 means managers have no influence on **non-oil profits**, 100 means a station's **non-oil profits** is solely determined by its manager.

____

Q30. In the past 6 months, how do you divide your working time into the following tasks now? Please indicate the time share (in %) of each task with a number between 0 and 100 and the total number should add to 100. [0-100]

- Strategic tasks (adjust the oil prices optimally, product mix and product price in the convenience store, seek new corporate customers) __________
- Analytical tasks (check and analyze data from the internal digital platforms and include the analysis in decisions) __________
- Operational tasks (assist employees with fueling, car washing, and customer service; Inspect gas stations, inspect convenience stores, post price tags, and ensure shelves are full; prepare for inspections from police stations, fire departments, and safety bureaus) __________
- Employee management (educate/train employees, talking to employees to motivate them, monitor employee to make sure they work hard, punish bad-performing employees, reward good-performing employees) __________

Q31. Please use a number between 1 and 5 to indicate whether you agree or disagree with the following statement, where 1 means you completely disagree and 5 means you completely agree. The larger the number, the more you agree. [1-5]

- Overall, I am satisfied with my work__________
- Work gives me a great sense of self-satisfaction__________
- I am satisfied with my salary__________
- I feel that my work is recognized by my superior__________
- I feel there are many opportunities for advancement__________

Q32. On a scale of 0 to 10, how much autonomy do you have over the daily operations of your station? [0-10]
0 means you have no control over the daily operations and all decisions are made by the higher level management, and 10 means you have full control of your station.

____

Q33. On a scale of 0 to 10, how much autonomy do you delegate to your employees?[0-10]
0 means you make every decision in your station and 10 means you fully delegate your autonomy to your employees and let them decide.

____

Q34.Are you satisfied with your current level of autonomy or would you like to have more or less decision rights?

- I'm satisfied__________
- I would like to have more autonomy__________
- I would like to have less autonomy__________

Q35. How well does each of the following statements describe you as a person? [1-5]
Please answer according to the following scale: 1 means "does not apply to me at all",5 means "applies to me perfectly." With values between 1 and 5, you can express where you lie between these two extremes.

- I can operate my station better with more autonomy__________
- I feel more valuable with more autonomy__________
- I feel more stressed with more autonomy__________

Q36. Have you done the following practices? [Please put in Yes/No/Not applicable]

- Visited competitor's stations to see prices in the **last month**__________
- Believed that the optimal price would be to match the competitor's price in the **last month**__________
- Gathered information about what products competitors have for sale, to help determine what products to sell in the **last month**__________
- Asked existing customers whether there are any other products the customers would like the station to sell or produce in the **last month**__________
- Talked with at least one former customer to find out why former customers have stopped buying from this station in the **last month**__________
- Asked a supplier/other station managers about which products are selling well in the convenience store in the **last month**__________
- Used a customized offer to attract new customers in the **last month**__________
- Advertised in any form in the **last month**__________ (if "Yes", please specify__________)
- Achieved having the stocks of oil and non-oil products not run out any more often than once per month __________
- Used records to see how much profits the station is making in the **last month**__________
- Used records regularly to know whether sales of a particular product are increasing or decreasing from one month to another in the **last month**__________
- Worked out the cost to the station of each main product it sells in the **last month**__________
- Know which goods you make the most profit per item selling in the **last month**__________
- Had a written budget, which states how much is owed each month for rent, electricity, equipment maintenance, transport, advertising, and other indirect costs to business in the **last month**__________
- Reviewed the financial performance of the business and analyzed where there are areas for improvement in the **last month**__________
- Communicated a clear understanding to employees of the expectations from their jobs in the **last month** __________
- Provided coaching and guidance on how employees can improve performance **in the last month**__________
- Actively supported employee professional/career development in the **last month**__________

- Consult with employees for decision-making when appropriate, **in the last month**__________
- Worked to generate a positive attitude in the team, even when conditions are difficult **in the last month**__________
- Did something **in the past month** to demonstrate to employees that you are someone they can trust__________
- Gave relatively higher bonuses to better performing employees in the **last month**. __________
- Gave relatively lower bonuses to worse performing employees in the **last month**. __________

Q37. What share of the employee's performance pay is subject to your adjustment? [0-100]
Please answer with a number between 0 and 100, where, for example, 10 means 10% of the performance pay.

____

Q38. How much more in percentage points does the most productive worker earn compared to the least productive worker in your station in a typical month? Please answer with a number between 0 and 100. [0-100]
Here 0 means the most productive worker earn just as much as the least productive worker, and 100 means the most productive worker earn 100% more than (or double of) the least productive worker.

____

Q39. In general, how willing or unwilling you are to take risks, using a scale from 0 to 10, where 0 means you are "completely unwilling to take risks" and 10 means you are "very willing to take risks." [0-10]

____

Q40. We now ask you for your willingness to act in a certain way. Please again indicate your answer on a scale from 1 to 5.
A 1 means "completely unwilling to do so," and a 5 means "very willing to do so."

- How willing are you to give up something that is beneficial for you today in order to benefit more from that in the future? __________
- How willing are you to punish someone who treats you unfairly, even if there may be costs for you? __________
- How willing are you to punish someone who treats others unfairly, even if there may be costs for you? __________

- How willing are you to give to good causes without expecting anything in return? __________

Q41. How well does each of the following statements describe you as a person? Please indicate your answer on a scale from 1 to 5.
A 1 means "does not describe me at all," and a 5 means "describes me perfectly."

- When someone does me a favor, I am willing to return it. __________
- If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so. ____________________
- I assume that people have only the best intentions. __________
- I tend to postpone tasks even if I know it would be better to do them right away. __________

Q42. Imagine the following situation: Today you unexpectedly received 800 [local currency]. How much of this amount would you donate to a good cause?

____

Q43. Please think about what you would do in the following situation. You are in an area you are not familiar with, and you realize that you lost your way. You ask a stranger for directions. The stranger offers to take you to your destination.
Helping you costs the stranger about 16 [local currency] in total. However, the stranger says he or she does not want any money from you. You have six presents with you. The cheapest present costs 4 [local currency], the most expensive one costs 24 [local currency]. Do you give one of the presents to the stranger as a "thank you" gift?
(If yes, ask:) Which present do you give to the stranger? (Read 2-7)

1. No, would not give present
2. The present worth 4 [local currency]
3. The present worth 8 [local currency]
4. The present worth 12 [local currency]
5. The present worth 16 [local currency]
6. The present worth 20 [local currency]
7. The present worth 24 [local currency]

Q44. Suppose that a few days ago two individuals, let us call them worker A and worker B, were recruited via an online market place to conduct an assignment. They were each offered a participation compensation of 12 [local currency] regardless of what they were paid for the assignment. After completing the assignment, they were told that their earnings from the assignment would be the same regardless of the number of tasks they completed. However, they were also told that a third person would be informed about the assignment and their productivity, and would be given the opportunity to redistribute the earnings and thus determine how much they were paid for the assignment.
You are the third person and we now want you to choose how earnings should be distributed between worker A and worker B. Your decision is completely anonymous. The workers will receive the payment that you choose for the assignment within a few days, but will not receive any further information. Worker A completed more tasks, while worker B completed less tasks. However, Worker A and worker B were each paid 18.0 [local currency] for the assignment. Please state whether and how you would like to redistribute between them.
Please state which of the following alternatives you choose:

1. **I do not redistribute:** worker A is paid 18 [local currency] and worker B is paid 18 [local currency].
2. **I do redistribute**
   (a) worker A is paid 36 [local currency] and worker B is paid 0 [local currency].
   (b) worker A is paid 30 [local currency] and worker B is paid 6 [local currency].
   (c) worker A is paid 24 [local currency] and worker B is paid 12 [local currency].
   (d) worker A is paid 12 [local currency] and worker B is paid 24 [local currency].
   (e) worker A is paid 6 [local currency] and worker B is paid 30 [local currency].
   (f) worker A is paid 0 [local currency] and worker B is paid 36 [local currency].

Q45. Suppose that a few days ago two other individuals, let us call them worker C and worker D, were recruited via an online market place to conduct an assignment. They were each offered a participation compensation of 12 [local currency] regardless of what they were paid for the assignment. After completing the assignment, they were told that their earnings from the assignment would be determined by a lottery. The worker winning the lottery would earn 36 [local currency] for the assignment and the other worker would earn nothing for the assignment. They were not informed about the outcome of the lottery. However, they were told that a third person would be informed about the assignment and the outcome of the lottery, and would be given

the opportunity to redistribute the earnings and thus determine how much they were paid for the assignment.
You are the third person and we now want you to choose how earnings should be distributed between worker C and worker D. Your decision is completely anonymous. The workers will receive the payment that you choose for the assignment within a few days, but will not receive any further information. Worker C won the lottery and earned 36 [local currency] for the assignment, thus worker D earned nothing for the assignment.
Please state which of the following alternatives you choose:

1. **I do not redistribute:** worker C is paid 36 [local currency] and worker D is paid 0 [local currency].
2. **I do redistribute**
   (a) worker C is paid 30 [local currency] and worker D is paid 6 [local currency].
   (b) worker C is paid 24 [local currency] and worker D is paid 12 [local currency].
   (c) worker C is paid 18 [local currency] and worker D is paid 18 [local currency].
   (d) worker C is paid 12 [local currency] and worker D is paid 24 [local currency].
   (e) worker C is paid 6 [local currency] and worker D is paid 30 [local currency].
   (f) worker C is paid 0 [local currency] and worker D is paid 36 [local currency].

Q46. Please answer according to the following scale: 1 means "does not apply to me at all", 5 means "applies to me perfectly." With values between 1 and 5, you can express where you lie between these two extremes. [1-5]
I see myself as someone who...

- Does a thorough job__________
- Is communicative, talkative__________
- I sometimes somewhat rude to others__________
- Is original, comes up with new ideas__________
- Worries a lot__________
- Has a forgiving nature__________
- Tends to be lazy__________
- Is outgoing, sociable__________

- Values artistic experiences__________
- Gets nervous easily__________
- Gets things done effectively and efficiently__________
- Is reserved__________
- Is considerate and kind to others__________
- Has an active imagination__________
- Is relaxed, handles stress well__________

Q47. How competitive do you consider yourself to be? Please choose a value on the scale below, where the value 0 means "not competitive at all" and the value 10 means "very competitive".

____

Q48. In general, are you a person who is confident that you can do better than others, or are you not that confident? [Not at all confident (0)-Very confident (10)]

____

Q49. Please answer according to the following scale: 1 means "does not apply to me at all", 5 means "applies to me perfectly." With values between 1 and 5, you can express where you lie between these two extremes.

- When I make plans, I am almost certain that I can make them work.__________
- Getting people to do the right things depends upon ability; luck has nothing to do with It.__________
- What happens to me is my own doing.__________
- Many of the unhappy things in people's lives are partly due to bad luck.__________
- Getting a good job depends mainly on being in the right place at the right time.__________
- Many times I feel that I have little influence over the things that happen to me.__________
- I have a hard time breaking bad habits.__________
- I get distracted easily.__________
- I say inappropriate things.__________
- I refuse things that are bad for me, even if they are fun.__________
- I'm good at resisting temptation.__________

- People would say that I have very strong self-discipline.__________
- Pleasure and fun sometimes keep me from getting work done.__________
- I do things that feel good in the moment but regret later on.__________
- Sometimes I can't stop myself from doing something, even if I know it is wrong.__________
- I often act without thinking through all the alternatives.__________

Q50. How do you see yourself: In general, are you a person who likes to delegate authority, or are you a person who likes to retain authority? 0=does not like authority at all, 10=likes authority very much.

Q51. In the next question you can choose either Box K or Box U. Both hold 100 balls which can either be purple or orange? For Box K, the exact mix of purple balls and orange balls is given below. Box U also holds purple and orange balls, but the mix is unknown. In other words, both boxes hold 100 balls with two different colors (purple and orange). The mix of purple and orange balls is known for Box K and unknown for Box U. One ball will be drawn at random from the box you choose. Imagine You will win $15 if a purple ball is drawn.

1. Box K
2. Indifferent
3. Box U

Q52. Imagine you are tossing a fair coin. After eight tosses you observe the following result: tails-tails-tails-heads-tails-heads-heads-heads. What is the probability, in percent, that the next toss is "tails"? Please answer with a number between 0 and 100.

____

Q53. Suppose you are matched with another station manager to play a game. Your opponent and you are going to ask for an amount of money from a referee for the game. The amount must be between $1 and $6. You will get the amount of money you ask for. However, you will get $10 more if you ask for exactly $1 less than your opponent. How much money do you ask for?

____

Q54. Considering all station managers who play this game, what percent do you think will earn less money than you? The same money as you? More money that you? (answers must sum to 100%)

- Less than you__________
- Same as you__________

- More than you__________

Q55-Q63. Raven's Progressive Matrices.

Q64-Q71. Reading the Mind in the Eyes Test.