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

Yi Han David Huffman Yiming Liu

March 7, 2026

Abstract

We investigate how bounded rationality of firm managers affects pricing strategies using data from a firm with over 20,000 gas station managers. Managers with lower cognitive skills exhibit systematically different mental models of optimal price, favoring lower prices, and this persists in the face of experience. This partly reflects not anticipating how competitors respond to price cuts. Consequences include more price wars, lower profits and producer surplus, but higher consumer surplus and market efficiency. Our findings demonstrate that cognitive constraints can systematically affect repeated price competitions, with important theoretical, welfare, and policy implications.

JEL classification: D22, D91, L1

Keywords: Bounded rationality; cognitive skills; mental models; narratives; level-k; cognitive hierarchy; endogenous depth of reasoning.

*Han: Renmin University of China; yihanecon@ruc.edu.cn. Huffman: Cornell University; dbh234@cornell.edu. Liu: Humboldt University of Berlin and WZB; yiming.liu@hu-berlin.de.

1 Introduction

Explaining prices is arguably one of the primary goals of economics. Mainstream theories of repeated price competition assume that firms are fully rational, and set prices in equilibrium, with the threat of retaliatory price cuts potentially allowing maintenance of high prices and profits. This abstracts from the reality, however, that firm managers are, to varying degrees, boundedly rational.[1] The effects of cognitive constraints are not obvious ex ante. There could be little effect, because high stakes and learning opportunities cause quick convergence to optimal pricing strategies. If constraints do matter, this could take the form of noise, with managers making errors centered around the optimal price. A more radical possibility would be if cognitive constraints lead managers to have systematically different ideas about what is the optimal pricing strategy. The distinction is important, for example because noise may have modest effects on market outcomes, averaging over time, whereas systematic differences in strategies could more substantially affect prices and market efficiency.

What might be a mechanism through which cognitive constraints foster systematically different ideas about optimal pricing? One candidate emerges from the branch of theory in behavioral economics that has considered strategic decision making, including level-k, cognitive hierarchy, and endogenous depth-of-reasoning (EDR) models (e.g., Nagel, 1995; Costa-Gomes et al., 2001; Camerer et al., 2004; Alaoui and Penta, 2016). A common theme is that bounded rationality leads to underestimating competitor sophistication, treating their behavior as non-strategic and exogenously determined. In the context of price competition, this could lead to a *systematic tendency to favor lower prices*, because of not expecting low prices to trigger competitor price cuts. Laboratory studies have linked a leading measure of cognitive skills, the Raven Matrix test, to ability to anticipate competitor behavior in "beauty contest" type games or repeated Prisoners' Dilemma games.[2] However, field evidence on how cognitive skills shape pricing has been lacking.

The goal of this paper is to provide a first, large-scale investigation of whether and how the cognitive constraints of firm managers affect pricing in repeated price competitions. We provide field evidence on three sets of open questions. Do lower cognitive skills systematically affect managers' mental models of optimal pricing strategy, and actual pricing decisions? If so, is this in the direction of favoring lower prices, as suggested by behavioral theories, and is there evidence that the specific mechanism of not anticipating competitor responses to price cuts plays a role? What are the consequences in terms of spillovers to market prices (price wars), profits, consumer surplus, market efficiency, and traditional measures

[1] We use the terms "bounded rationality," "cognitive constraints," and "low cognitive skills," interchangeably.

[2] See Gill and Prowse (2016), Proto et al., (2019), Fe et al., (2022), Gill and Rosokha (2026). Burnham et al. (2009) and Carpenter et al. (2013) provide related evidence.

of market power?

In brief, our main findings are as follows. (1) *Lower cognitive skills cause a systematic difference in beliefs about optimal pricing,* towards favoring lower prices, a self-reported desire to engage in price cuts, and setting substantially lower actual prices; (2) *Lower cognitive skills lead to not anticipating competitor sophistication and price responses,* and this partly mediates the link between cognitive skills and mental models of optimal price and actual pricing decisions; (3) *The lower prices arising from low cognitive skills have important consequences,* including spillovers to market prices by triggering competitor price cuts and price wars, reduction of firm profits and producer surplus (PS), higher consumer surplus (CS) and market efficiency, and biases in traditional approaches to measuring market power.

Answering our research questions is demanding in terms of data. It requires pricing decisions and profits for a large sample of managers, combined with measures that to our knowledge have never been implemented together in a field setting: manager cognitive skills, measures of mental models of optimal pricing, measures of mental models about competitor sophistication and price responses, and measures of other characteristics that should be controlled for such as discount factors and risk preferences. Welfare analysis also requires knowing marginal costs, since the typical approach of assuming rationality to infer marginal cost is problematic when bounded rationality is present. We obtain data with all of these ingredients by collaborating with a company operating more than 20,000 gas stations.

Our findings generate important implications for firms, policymakers, and economic theory. First, they provide an "existence proof" that cognitive constraints can affect pricing in real-world repeated price competition—even at large companies, where pricing is delegated to managers—with spillovers to market-level prices, and these effects persist despite experience and feedback. Second, bounded rationality systematically shifts managers toward favoring lower prices, which tend to reduce firm profits but increase CS and market efficiency. This implies that firms seeking to improve pricing may need to reshape managers' mental models rather than simply provide better data. For policymakers, these findings suggest that the increasing trend of replacing human price-setting with AI algorithms could substantially affect market prices and efficiency. Third, the underlying mechanism—failing to anticipate competitor price responses—could be quite general, creating a force for lower prices across a wide range of industries characterized by repeated price competition. This mechanism also adds nuance to mainstream theoretical explanations for fundamental economic phenomena such as price cuts and price wars, with important implications for competition policy: some price wars are likely strategic mistakes rather than telltale signs of optimal strategies to encourage competitors to keep prices high.

Section 2 details the nature of station manager jobs and the market environment, describes our data sources, and introduces a simple conceptual framework. Key aspects of the work setting include the fact that station managers are company employees with performance pay linked to station profits, and they can substantially influence fuel pricing, though

the company has some safeguards in place. Upper management explains giving pricing discretion as a trade-off: it can lead to mistakes but also harnesses local knowledge.

The market environment involves two large companies, including our partner, and many smaller, independent companies. There are government-imposed price ceilings for fuel products (gasoline and diesel), indexed to the world price of oil. These potentially facilitate coordination on high prices, and indeed, the default policy of the two large companies is to price at the ceilings, while independents usually track them with a fixed discount. This institutional feature provides a natural benchmark for identifying price wars. Additionally, because our partner company's internal accounting links changes in marginal cost one-to-one to ceiling movements, variation in the ceiling serves as a cost shifter for estimating demand, which underpins our welfare calculations. Factors that could affect pricing, such as station type, location characteristics, or number of local competitors, are quite balanced with respect to manager cognitive skills, mitigating concerns about omitted variable bias.

We use four data sources: (1) a survey of approximately 350 district-level senior managers on discretion given to station managers and their views on potential mistakes; (2) multiple survey waves with 20,000 station managers, achieving roughly 14,000 responses each time, measuring cognitive skills (including a Raven test), mental models of optimal pricing and competitor behavior, and a wide range of other traits we denote noncognitive skills (e.g., discount factor, risk aversion, personality) to serve as controls; (3) four years of monthly panel data on prices and profits for nearly all company stations; (4) daily pricing data for one region, enabling identification of price wars and demand estimation.

Section 2 concludes with a conceptual framework. Building on behavioral theory, the core logic of our framework is that low cognitive skills cause managers to not anticipate that competitors respond strategically to price cuts, which makes cutting price appear more profitable and leads to choosing systematically lower prices. We use the framework to guide our investigation of underlying mechanisms: It implies that mental models of competitor behavior should mediate the link between cognitive skills and pricing, though we do not claim this is the only channel through which cognitive skills could systematically influence pricing. We illustrate how the framework can be made more formal using a level-k model: Low-skill managers are level-1 types who view competitors as non-strategic level-0 types; they seek to undercut competitors to steal market share, potentially generating price wars when playing against other level-1 types. Higher-skill level-2 types believe competitors are level-1 types who respond to price cuts with further cuts, and thus find it profitable to maintain higher prices when patient enough. Neither Level-1 nor level-2 adjust their beliefs about opponent types, because they attribute unexpectedly low or high competitor prices to unobserved cost shocks. Although our application is to retail fuel markets, the logic applies broadly across price competition settings.

Section 3 presents our first set of results on how cognitive skills affect mental models of optimal pricing and actual pricing decisions. We asked managers an open-ended ques-

tion about what contributes most to consistently high fuel profits, and find systematically different ideas about optimal pricing: High-skill managers were significantly more likely to emphasize maintaining high prices (e.g., "Do not blindly engage in price wars"), while low-skill managers stressed high volume and low prices (e.g., "Increase sales through price cuts"). Lower cognitive skills also significantly predict a self-reported desire to cut prices. Turning to actual prices, managers with lower cognitive skills charge significantly lower prices on average, controlling for noncognitive skills, station, location, and market characteristics. The magnitude is substantial: a one-standard-deviation reduction in cognitive skills is comparable to adding three-fifths of a competitor to the local market. To provide an even more conservative test of causality and shed light on how the impact evolves over time, we conducted an event study, defining treated stations as those receiving a new manager, with control stations having no manager change. This holds constant everything about the station and market. Over the two-year post-period, low-skill managers reduce prices whereas high-skill managers raise them. This indicates that experience is not a substitute for high cognitive skills in our setting, and suggests that managers with different cognitive skills are "learning" different things from the same environment, potentially by viewing their experience through the lenses of different mental models of competitor behavior.

Section 4 investigates whether or not anticipating competitor strategic behavior could be a mechanism. As one measure of ability to anticipate competitors, we implemented a hypothetical version of the money-request game (Arad and Rubinstein, 2011; Fe et al., 2022). Managers could request an amount from \$1 to \$6, receiving whatever they requested, but earning an additional \$10 if they requested exactly \$1 less than their opponent. A key feature of this game is a salient, intuitive response for non-strategic players: Requesting \$6. We observe a modal choice of \$5, consistent with many managers assuming non-strategic opponents, and these managers have relatively low cognitive skills. With so many managers requesting \$5, the ex-post optimal choice was \$4. Managers with high cognitive skills were significantly more likely to choose \$4. The choice of \$4 among high-skill managers provides a model-free indicator of the ability to anticipate strategic behavior, and the link to high cognitive skills is consistent with our conceptual framework. We also designed two survey questions about how managers think about competitors in fuel markets. One shows that lower-skill managers are *more* confident about their ability to influence fuel sales, consistent with believing that prices can be lowered without triggering competitor price responses that would dampen sales. The other shows that lower cognitive skills predict believing it is optimal to cut price when being undercut, consistent with not anticipating further reductions by competitors.

We next provide more direct evidence that our measures of cognitive skills and mental models of competitors are capturing a tendency to anticipate competitor price responses. We showed a subsample of managers the data from a real station in which a competitor responded to our partner station's price cut with a deeper cut, and asked for counterfactual

predictions about sales and competitor prices if our station had cut its price even earlier. Low-skill managers were significantly more likely to predict no competitor response and, conditional on predicting a change, were as likely to predict an increase as a decrease—consistent with treating competitor prices as exogenously determined. By contrast, high-skill managers were significantly more likely to predict negative price responses. Our three measures of ability to anticipate competitor behavior—requesting $4 and the two survey measures—are also significantly associated with predicting competitor price cuts in the expected directions. We conclude with a mediation analysis, showing that the measures of anticipating competitor behavior help mediate the relationship between cognitive skills and narratives about fuel profits, desire to cut prices, and actual prices. We also check robustness to assuming a specific model of bounded rationality, level-k; through this lens, the money request game identifies manager level of reasoning. Consistent with our level-k model, cognitive skills explain whether managers are level-1 (requesting $5) or level-2 (requesting $4), level-1 managers are significantly less likely to anticipate competitor price responses, and they charge lower fuel prices.

Section 5 analyzes economic consequences and implications for measuring market power. We start with whether the price cuts of low-skill managers spill over to market prices, and find that low-skill managers are involved in about twice as many price wars as high-skill managers. Turning to whether the low price strategy is biased downwards compared to what is optimal for the firm, we exploit exogenous cost-shifters to estimate station fuel demand, allowing inference about the effect of price cuts on producer surplus. We find that the price decrease associated with a one-standard-deviation decrease in cognitive skills reduces profit by 2.6%. On the welfare side, it is clear that lower prices set by low-skill managers must improve CS and market efficiency, under weak assumptions of downward-sloping demand and market power; we use our demand estimates to quantify these effects, finding that a one-standard-deviation decrease in cognitive skills reduces deadweight loss (DWL) by 6.2%. Conversely, increasing cognitive skills—mimicking the replacement of human price-setters with sophisticated AI algorithms—would substantially increase DWL. Finally, we discuss how standard approaches to measuring market power, based on assuming rational pricing, will be biased and will misattribute price variation driven by cognitive skills to other factors, such as cost differences.

Section 6 concludes. We summarize key findings and some further implications, briefly discuss results regarding noncognitive skills and experience (which have little impact on pricing decisions), and point to directions for future research. For example, in follow-up work we are studying whether some managers may have a flawed understanding of consumer responses to price changes.

Our paper contributes to a nascent literature documenting mistakes by "behavioral firms" (for a survey see Heidhues and Kőszegi, 2018).[3] We are novel in directly measuring man-

[3]Examples include uniform pricing by retail chains (DellaVigna and Gentzkow, 2019), suboptimal adver-

agers' mental models of optimal strategies, showing that a measure of cognitive skills can explain systematic differences in both these mental models and pricing decisions, and identify a channel through which bounded rationality of firm managers may increase market efficiency.[4] We also complement previous field studies that have used level-k-type models to explain suboptimal firm behavior.[5] We are different in studying the setting of repeated price competitions, and showing how bounded rationality leads to systematically lower prices, with benefits for efficiency. Our approach is also distinct: rather than inferring strategic types from field behaviors, we use direct measures of beliefs about competitor strategic behavior and price responses—key mechanisms in level-k and related models—and show that these are linked to underlying differences in cognitive skills, and have explanatory power for both perceived optimal prices and actual pricing decisions.

We also contribute to two other behavioral literatures: on strategic competition in the lab and on narratives and mental models. Our findings complement lab studies of beliefs and strategic heterogeneity in repeated games (e.g., Dal Bó and Fréchette, 2019; Aoyagi et al., 2024; Gill and Rosokha, 2024 and 2026) by documenting heterogeneity in real strategic competitions and linking this to cognitive skills and beliefs about competitor sophistication and price responses.[6] While lab studies have shown level of reasoning increasing with experience in tightly-controlled environments that rule out factors like exogenous shocks (e.g., Duffy and Nagel, 1997; Danz et al., 2012; Gill and Rosokha, 2026), our findings show that experience can be less effective in more complex field environments (our level-k model incorporates this complexity by assuming that managers can mis-attribute competitor price movements to cost shocks rather than revising beliefs about opponent types). Turning to the second literature, we contribute to growing work on misspecified mental models and narratives in economic decisions (e.g., Kendall and Charles, 2022; Montiel Olea et al., 2022; Andre et al., 2023a; Andre et al., 2023b; Esponda et al., 2024), offering new evidence that differences in mental models can be explained by cognitive skills, that these differences

tising by inexperienced online sellers (Tadelis et al., 2023), failure to exploit consumer left-digit bias (Strulov-Shlain, 2023; List et al., 2023), and overreaction to transitory shocks by less experienced bar owners (Goldfarb and Xiao, 2024). There is also an older tradition of bounded rationality in IO, which assumes firms use rules of thumb or face explicit cognitive costs (for a survey see Ellison, 2006).

[4]We do not claim boundedly-rational managers necessarily increase overall efficiency, because they could reduce efficiency in other ways, e.g., misallocation of workers to tasks. The tendency to charge lower prices implies a countervailing effect, however, and policies that govern pricing, such as sophisticated AI algorithms or centralized pricing, may harm efficiency.

[5]Goldfarb and Yang (2009) study internet service providers' technology adoption decisions, using competition avoidance to identify strategic sophistication, and find that such firms have higher survival rates. Goldfarb and Xiao (2011) show that manager education, proxying for sophistication, predicts entry into less competitive telephone markets and higher survival rates. Hortaçsu et al. (2019) demonstrate that a cognitive hierarchy model nicely organizes electricity firm bidding, with less sophisticated types—typically smaller firms—identified by having steeper bid functions than predicted by a rational auction model. Because high bids can price low-cost firms out of the market, this can reduce productive efficiency.

[6]Our findings are also consistent with lab evidence that implementing certain strategies may be cognitively demanding (Oprea, 2020; Proto et al., 2022).

influence real economic outcomes, and that they persist despite high stakes and feedback.[7]

Our findings also speak to mainstream literatures on strategic price competition. They add nuance to the view that price wars signal collusion (e.g., Green and Porter, 1984; Slade, 1992), showing that some may instead result from cognitive mistakes. By revealing that markups may be distorted by manager cognitive skills, they contribute to the literature on measuring market power (see Berry et al., 2019). They also add to research on whether price ceilings serve as focal points (e.g., Knittel and Stango, 2003), showing that such coordination may depend on managers' cognitive skills. Our results complement evidence that algorithmic pricing can increase market prices—e.g., Assad et al. (2024) find that algorithmic pricing raises fuel prices in Germany—by suggesting a reason: human pricing may leave market power partially unexploited.[8] Our results also complement previous work using retail gas markets to test models of strategic pricing (e.g., Hastings, 2004; Noel, 2007; Barron et al., 2008; Houde, 2012; Luco, 2019) by incorporating the role of managerial cognition.

A related literature in economics has demonstrated that managers matter in various ways for 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 people-management roles, we study how managers matter for firms through their pricing decisions. Our paper also adds to 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), showing a non-obvious impact of low cognitive skills on pricing decisions, towards systematically lower prices, which implies a way that low cognitive skills can improve efficiency.

2 Market setting and Data

2.1 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 nearly always have a convenience store. In this setting, the bulk of station profits come from fuel sales rather than convenience stores; on average, fuel profits make up 71% of total station profits. The stations are primarily company-owned rather than franchises. Each station has a manager with substantial influence over operations, including pricing decisions. Station operations are also governed by district-level managers, who set policies about the type and degree of discretion given to station managers.

[7] 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).

[8] By shedding new light on the properties of human price setting, our findings also complement a growing literature using simulations and laboratory studies to investigate algorithmic pricing (e.g., Calvano et al., 2020; Hansen et al., 2021; Asker et al., 2024; Fish et al., 2024; Arunachaleswaran et al., 2024).

There are approximately 350 districts, each overseen by a district manager.

The market is relatively concentrated with significant brand differentiation. Our partner company is one of two major brands, each claiming about one-third of the market, with stations nationwide. The remainder goes to hundreds of smaller, local chains, which we refer to as "independent" competitors. A key difference is vertical integration: major brands produce their own fuel, while independents purchase from local refineries on the open market. Major brands also position themselves as offering premium products and typically charge more than independents for the same grade of fuel.

The government imposes a separate price ceiling for each fuel product—primarily gasoline grades (92, 95 and 98-octane) and diesel—each indexed to the world price of oil and adjusted every 10 working days based on cumulative changes in that index. The ceilings move in parallel: a given change in world oil prices produces the same absolute adjustment across products, though the level of each ceiling differs to reflect product-specific refining costs. Because our partner company is vertically integrated, its internal transfer prices from the upstream refining branch move one-to-one with the ceilings, providing a direct measure of marginal cost variation that becomes important for estimating demand elasticity in Section 5. Independent competitors, by contrast, purchase from local refineries and face different cost dynamics. The ceilings also serve as natural focal points for price coordination: the two major brands price near these most of the time, while independents typically track ceiling movements, although at a substantial discount.

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 station manager at the current assignment. "Competitors in local market" refers to stations from other companies within 2.5 km, though the definition also accounts for 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.

Table 1 shows descriptive statistics for station managers and their stations. The median manager age is 39, and about 70 percent are male. The modal education level is a junior-college degree. Managers stay with the company for substantial periods, with median company tenure of 7 years, though they switch stations periodically, with median station tenure of about 2 years. The median number of employees is 5, so managers have people management duties but not for large teams. The median number of competitors in the local market

is 2. The company defines local markets starting with a 2.5 km radius, but then adjusts for factors like commuting routes. The median market share for a partner company station is about 30 percent of local fuel sales.

2.2 Study timeline and rationale behind survey design

Our collaboration with the partner firm began in 2021, through a contact in the company's research arm.[9] The company requested evaluation of a policy and, in exchange, allowed us to survey station managers and provided data on pricing and station performance. We judged that the most unique aspect of the opportunity was access to a large sample of managers with influence over strategic decisions, including fuel pricing. The behavioral literature on strategic decision-making, most notably level-k, cognitive hierarchy, and EDR models, had focused primarily on relaxing the assumption of full rationality. We therefore chose to explore how bounded rationality might influence pricing decisions, and specifically to test the mechanism these models hypothesize: not anticipating competitor strategic behavior.

Initial interviews with station managers provided anecdotal evidence about strategic thinking. Some managers, whom we understood to be relatively successful, emphasized keeping fuel prices high and avoiding price wars. Others expressed frustration about senior management reluctance to let them cut prices, proactively to increase sales, or in response to being undercut; their thinking appeared to treat competitor prices as exogenously determined. One manager described how a price cut "worked at first, but then stopped working," without a clear explanation—suggesting a blind spot regarding competitor price responses. These interviews highlighted the value of measuring cognitive skills and anticipation of competitor behavior as potential explanations for heterogeneity in pricing behavior.

We then designed surveys administered through the research arm's internal infrastructure. Because participation is expected, response rates are high; because surveys are administered independently of station managers' supervisors, assurances of confidentiality are credible. Since HR data were unavailable, our surveys also collected demographic variables.

We planned two survey waves with station managers and one survey with district-level managers. The core of each station manager wave was a test of cognitive ability and the money request game—a strategic game from the level-k literature—to measure ability to anticipate competitor behavior. Surveys also measured other manager traits, such as economic preferences and personality type, to serve as controls. We included these same trait measures in both waves to assess measurement error and to ensure contemporaneous measures for explaining wave-specific outcomes. Survey length constraints required us to spread additional measures across the two waves. The first wave included a management practices inventory with questions probing managers' understanding of competitor price responses. The second wave included an open-ended "narratives" question eliciting manager beliefs

[9] IRB approval was granted by Renmin University of China.

about determinants of fuel profits, including views on optimal pricing. The district manager survey asked about the authority given to station managers, reasons for delegation, and upper management concerns about strategic mistakes. We launched the first manager survey and the district manager survey in 2022, and launched the second manager survey in 2023.

In 2024, we were given a chance to conduct a third wave. We primarily designed this wave for other projects but included a new measure of whether station managers anticipate competitor price responses. The measure presented real data from one of our partner stations cutting its price and a competitor responding with a deeper cut. We then elicited counterfactual predictions about what would have happened if the station had cut its price even earlier, testing whether managers predicted a negative price response. The wave again measured cognitive ability and included the money request game.

Each station manager wave was sent to all 20,000 managers. Response rates were consistently around 70 percent, yielding approximately 14,000 responses per wave. We also received responses from 353 district managers, a nearly 100 percent response rate. For reader convenience, when reporting results based on survey measures we give question wording in the text or footnotes, but survey instruments are also in Appendices F, G, H, and I.

2.3 Datasets on station performance and pricing decisions

The company provided monthly performance data for its gas stations from 2019 to 2022. These panel data record fuel and nonfuel profits, sales volume in gallons, and average monthly prices for gas and diesel products. We were given access to data from 26 of the company's 31 regions.[10] The dataset covers about 16,000 gas stations. The company matched performance data to first-wave survey responses using an internal identifier, yielding linked data for roughly 10,000 stations (reflecting the 70 percent response rate). For the second and third waves, the company did not perform this matching, but we linked responses across waves using other survey identifiers. This allows linking second- and third-wave responses to performance data, but only for managers who also responded to the first wave.

Another dataset provides daily price and sales data for one region, including prices and characteristics of all competitors. This region has roughly 1,300 partner company stations, with daily panel data from 2018 to 2021. These data allow us to identify price wars using competitor price information and to estimate demand parameters for welfare analysis. Appendix Table A.1 summarizes our data sources, indicating which measures are included in each survey wave and in the administrative performance data.

[10]Of the remaining five, some record data only quarterly and others were unavailable for administrative reasons.

2.4 Managerial discretion over fuel prices

Station managers in our partner company have substantial influence on fuel prices. In our survey of district managers, 48% of districts report that station managers can change posted prices without submitting a proposal (most often within a pre-specified range).[11] In the remaining districts where proposals are required, the average approval rate is about 34%.[12] Thus, managers can influence prices through multiple channels: directly adjusting prices, or submitting proposals that are approved about one-third of the time.

Why does the company allow such discretion? Our survey of district managers points to two main reasons. First, station managers can assess local factors difficult for upper management to observe, such as how responsive local competitors are likely to be to price changes.[13] Second, there is a strategic deterrence rationale: as one district manager noted, if station managers lack flexibility to respond quickly, competitors can easily undercut them.[14] The ability to react potentially provides a credible threat that helps sustain higher prices.

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 two KPIs: fuel profits and nonfuel (convenience store) profits. Performance on each KPI is assessed relative to a target, multiplied by a weight, and summed to determine bonus pay (Appendix A.2 provides more details). The average weight on fuel profit exceeds that on nonfuel profit, and targets are calibrated to closely match historical profits.[15] The most relevant feature for our analysis is that fuel profits directly enter the bonus formula, and managers have incentives to set fuel prices to maximize fuel profits.[16] Managers may differ, however, in their beliefs about the

[11]The question asked: "Please indicate the extent to which station managers must submit proposals for the following business decisions. Please answer each item with a number from 1 to 4. 1=Full autonomy, no reporting or proposal required; 2=Can implement what they want but must report; 3=Can choose freely within a pre-specified range, approval required beyond this range; 4=All decisions must seek approval." For "Adjusting fuel prices through direct price reductions," 48% chose a response other than 4.

[12]We asked, for different business decisions: "On average, what is the approval rate for proposals submitted by station managers?" The average reported rate for "Adjusting fuel prices through direct price reductions" was 34%.

[13]We asked "In your view, what are the characteristics of stations that are more suitable for granting more autonomy?" District managers could indicate all factors that applied from a multiple-choice list, and the most commonly chosen item was "Stations where the manager's knowledge of the local market is more important for profitability (e.g., stations facing more competition)."

[14]This response came from an open-ended question: "Do you think it would be a good idea for your gas stations to participate more actively in price competition and charge lower prices than competitors? Why or why not?"

[15]Incentive scheme weights are always the same for stations in a given district, and our analysis of pricing decisions controls for any differences in weighting schemes across districts using district fixed effects. Such multi-component KPI structures are common in managerial compensation (Kaplan and Norton, 1992).

[16]In our setting, managers do not have an incentive to pursue a loss-leader strategy using fuel prices—reducing fuel prices below the fuel-profit-maximizing level, hoping that increased nonfuel profits would more than compensate. Fuel is the main profit source for stations (more than 70% of total profits), and reducing

price level that achieves this goal, potentially depending on cognitive skills.

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

This section describes the manager traits we measured in our survey waves, 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 summarizes all traits, beginning with cognitive ability.

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 distribution (-)
Personality type	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 test (+)
Gender	Female indicator
Age	In years
Experience	In months

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 loadings on the corresponding factor are shown in parentheses (first survey wave).

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—the ability to generalize from an observed pattern to a general rule—which is a key element of fluid intelligence (Carpenter et al., 1990).[17] 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 A.1 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

price below the fuel-profit-maximizing level can be very costly. Indeed, the correlation of fuel price with total profits is positive (0.082), reflecting a positive correlation with fuel profits (0.091) and only a weak negative correlation with nonfuel profits (−0.017).

[17]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."

is consistent with the test measuring a trait of bounded rationality that is broadly relevant to strategic decision-making. For our particular application—strategic price competition in fuel markets—it seems plausible that inductive reasoning ability is a crucial determinant of station managers' ability to form accurate mental models of their competitors. Beyond capturing depth of reasoning, it may capture ability to generalize from the data they observe in their markets to form accurate models of the 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.[18] 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, including 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 preference, time preference, altruism, positive reciprocity, negative reciprocity, and trust—were measured using the survey module from the Global Preference Survey (Falk et al., 2018). These survey measures include quantitative items that mimic incentivized laboratory experiments, in addition to qualitative self-assessments, and 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). To measure "emotional intelligence" we use an eight item test, showing photographs of a person's eyes and asking respondents to identify the emotion expressed by the person's facial expression (Baron-Cohen et al., 2001).[19]

Factor analysis: To reduce dimensionality and capture latent traits underlying groups of measures, we performed factor analysis. More details are provided in Appendix A.4. 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 the first and second survey waves (Tables A.2 and A.3). 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

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

[19]This is an aspect of intelligence that is sometimes also denoted "theory of mind," and can be considered part of cognitive skills more broadly defined (Gill et al., 2025)

loadings quite similar across waves (Table A.4). 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.5). This factor loads positively on traits like conscientiousness, agreeableness, locus of control, confidence, patience, and emotional intelligence, 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 authority, but its eigenvalue is near 1 or below. 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.[20] Results of these checks are available upon request, but we provide an example in Figure B.6, where we show that results on the relationship of cognitive skills to fuel pricing decisions are robust to controlling for the noncognitive traits separately. 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 of such a bias in cognitive skill coefficients; instead, correcting for measurement error shows a stronger effect of cognitive skills.[21]

2.7 Balance of station, location, and market conditions by cognitive skills

Based on conversations with managers and company contacts, we know that the process of assigning managers to stations over time is complicated, somewhat arbitrary, and requires

[20]We check this because emotional intelligence could be considered part of cognitive skills, broadly defined, and a subset of items of the measure loaded on the same factor as the Raven test and numeracy measure in the second wave, although in the first wave these loaded with the other noncognitive skills measures.

[21]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 coefficient estimates we report for cognitive and noncognitive skills in our main analysis are lower bounds. Appendix A.5 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).

substantial lead-time to arrange. Managers also have some scope to request moves, raising the possibility of self-selection. Given our goal of testing whether cognitive skills affect pricing, it is of interest whether managers with different cognitive skills end up assigned to systematically different types of station, location, or market conditions, which could have an independent influence on pricing. For example, if managers with lower cognitive skills face more competition, this could lead them to charge lower prices. Table A.6 shows that these attributes are quite balanced across managers with high and low cognitive skills, with differences being economically small. This helps mitigate omitted variable concerns, but we also test robustness to controlling for these observable factors in regressions, and to using an event-study design that holds constant everything about the station and local market, including unobservables.

2.8 Conceptual framework

We develop a conceptual framework for one way that bounded rationality can matter in repeated price competitions. We build on a common theme across behavioral models of strategic interaction—including level-k, cognitive hierarchy, and endogenous depth of reasoning (EDR)—that bounded rationality leads to treating competitor behavior as non-strategic and exogenously determined.

We propose the following simple framework: (1) Low cognitive skills cause a manager to not anticipate that competitors will respond strategically to one's own pricing decision, and in particular, that price cuts trigger competitor price cuts; (2) because of not expecting competitors to cut prices in response, managers with low cognitive skills will perceive greater profitability of cutting price, making them more likely to choose a lower price.

The framework guides our investigation of the underlying mechanisms. Specifically, it implies that our measure of cognitive skills should predict, in a particular direction, both beliefs that price cuts are profitable and actual pricing decisions. Furthermore, cognitive skills should predict the ability to anticipate competitor behavior and price responses, as measured by our lab-in-the-field and survey-based measures of mental models of competitors. Finally, the measures of how managers think about competitors should at least partially mediate the link between cognitive skills and pricing beliefs and behavior. If mediation is only partial, this could reflect error in our measures of mental models of competitors, or the presence of additional mechanisms; we do not claim that the mechanism we test is the only channel through which cognitive skills could systematically influence beliefs and decisions about pricing.

While our framework does not take a stand on a particular way of modeling boundedly-rational reasoning, we illustrate how the framework can be made more formal by developing a level-k example (details in Appendix A.7.1). We focus on level-k for concreteness but view cognitive hierarchy and EDR models as plausible alternatives that generate similar

predictions. We consider a two-period Hotelling model in which level-1 firms—the more boundedly rational type—believe their competitors are non-strategic level-0 types who price as monopolists (or at the price ceiling if it is binding). Perceiving no strategic price response, level-1 firms undercut to steal market share; this can generate price wars when both firms are level-1 types. Level-1 types in our model do not question their belief that competitors are level-0, because they attribute unexpectedly low competitor prices to unobserved lower marginal costs (or equivalently, unobserved demand shocks). Level-2 firms believe their competitors are all level-1 types and thus anticipate that competitors observe and respond to prices, leading them to maintain higher prices to avoid triggering aggressive responses. We show theoretically that level-2 firms charge higher prices than level-1 firms given sufficient patience (level-2 types also maintain their beliefs over time, mis-attributing unexpectedly high prices to cost shocks). This framework maps naturally to cognitive skills as measured by the Raven test, which captures inductive reasoning—the ability to generalize from one's own strategic thinking to anticipate similar reasoning by competitors. We thus conceptualize lower cognitive skills as leading managers to be level-1 rather than level-2 types. While our empirical approach is primarily model-free, we also check robustness to using choices in a lab-in-the-field game to directly identify level-1 and level-2 types and test whether these are related to cognitive skills and beliefs and decisions about pricing as predicted by the model.

The next section tests whether lower cognitive skills lead to systematically different beliefs about optimal pricing and different price choices. The subsequent section narrows the focus to testing the underlying mechanism—anticipating competitor responses—highlighted in our conceptual framework.

3 Impact of cognitive skills on mental models of optimal price and actual pricing decisions

To test whether cognitive skills systematically influence managers' mental models of optimal pricing and affect actual pricing decisions, we proceed in three steps. First, we examine whether managers with different cognitive skills hold different views on profit-maximizing pricing. Second, we analyze whether they differ in stated preferences for price cuts. Third, we examine actual pricing decisions.

3.1 Impact of cognitive skills on narratives about fuel profits

We used a "narratives" approach to assess mental models of fuel profit determinants and the role of pricing (Andre et al., 2023).[22] In our second survey wave, we asked: "Some managers

[22]Andre et al. (2023) define a narrative as an explanation of a past economic event—in our case, a manager achieving consistently high fuel profits.

consistently have higher fuel profits than other managers. What do you think are the most important practices that enable them to achieve this?" This "hands-off" approach reveals whether managers focus on pricing when considering key determinants of profits, without leading them to think this way. We complement this with a more structured question, asking directly whether managers favor high or low prices, discussed in the next section.

As detailed in Appendix B.1, most narratives mapped onto the profit formula (profit = margin × sales volume), with managers emphasizing either margins (e.g., maintaining high prices) or volume (e.g., boosting sales through low prices). Some responses, such as effort or ability, did not fit cleanly into either category. Based on examining a random sub-sample we developed a classification rubric with 15 categories (Appendix E). Each of the more than 15,000 responses was independently classified by two RAs, blind to our hypotheses, using this rubric, achieving 75 percent inter-coder agreement. We report results based on the authors reconciling disagreements (because we believe this minimizes measurement error), but results are very robust to alternative approaches (see below). In the classification, most managers (80%) mentioned only one narrative, some (17%) mentioned two, and very few mentioned more.

Figure 1 presents the frequency of different narrative types and their relationship to cognitive skills. Panel (A) shows that 25% of narratives cite margin-related causes, 45% mention sales volume, and 30% cite other causes or "don't know." Panel (B) ranks narratives by the average cognitive skills of managers who mention them: those citing margin-related narratives—most frequently high prices—tend to have higher cognitive skills, while those focusing on sales volume or low prices tend to have lower skills. To give a better sense of how cognitive skills are related to narratives that are not empirically negligible, Panel (C) excludes infrequent narratives. This shows that cognitive skills are highest among managers citing high prices, followed by those citing sales volume or low prices, and lowest among those citing ability, effort, or "don't know." Comparing between managers mentioning high prices and those mentioning narratives related to low prices (low prices, sales volume, or customer development, respectively), non-overlapping 95% confidence intervals show that cognitive skills are significantly higher for those mentioning high prices.[23]

A natural question is whether these differences are robust to controlling for other manager traits or for station and market conditions, which could affect beliefs and be correlated with cognitive skills (although station and market conditions are quite balanced, as discussed in Section 2.7). Probit regression results (regression "All" in Figure 2) show that

[23]Throughout our analysis, significance tests will mainly be shown via 95% confidence intervals, and unless stated otherwise, significance is at the five-percent level or better. Managers who mention high sales being due to location, which is not necessarily a statement about the benefits of charging a low price but rather may be about demand conditions, also have lower cognitive skills than those mentioning high price, but the difference is not significant. A caveat when reading Figure 1 is that some managers mention multiple narratives, so there is a potential correlation of error terms across observations that is not accounted for in the depicted confidence intervals. Our significance statements are therefore based on the confidence intervals shown in Figure B.4, which excludes managers who mention multiple causes to ensure independence of observations.

Figure 1: Narrative measure of mental models for high fuel profits

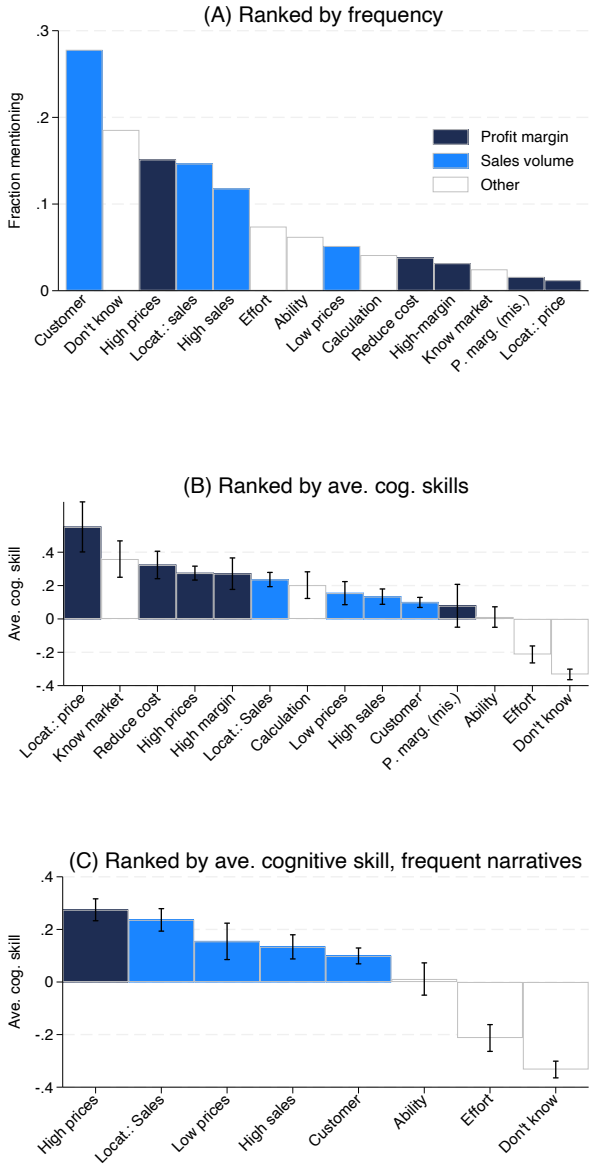

Notes: Panel (A) shows the frequency of narratives mentioning different categories of causes of high fuel profits. Bars are color-coded according to whether they fall into the profit margin, sales volume, or other categories. The narratives include: *Customer* (good customer service); *Don't know* (no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

the probability of mentioning high prices increases significantly with cognitive skills when controlling for noncognitive skills, other manager traits (age, gender, experience), and market and station characteristics. This partly reflects a tendency for higher-skill managers to

avoid non-price narratives included in the omitted category such as "don't know."[24] To focus more narrowly on whether cognitive skills predict different views on optimal pricing, the remaining regressions progressively tighten the omitted category to exclude narratives less obviously related to pricing. The most restrictive specification considers only managers who mention either high or low prices (regression "Low prices"). As shown in Figure 2, across all specifications higher-skill managers are significantly more likely to mention high prices, while lower-skill managers favor narratives about sales volume and low prices. Thus, two distinct mental models of optimal pricing emerge from the data, and these are strongly related to cognitive skills, with higher-skill managers emphasizing the importance of high prices and lower-skill managers emphasizing high sales-volume and low-price strategies.

Figure 2: Probability of mentioning high prices narrative by cognitive skills

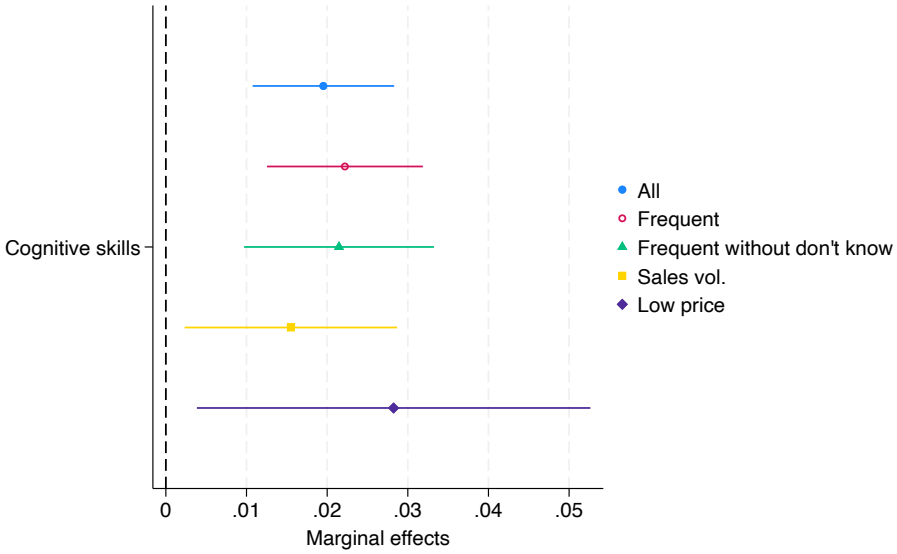

Notes: The figure reports marginal effects from Probit regressions with 95% confidence intervals. Each bar represents a separate regression showing how a one standard deviation increase in cognitive skills affects the probability of mentioning the high prices narrative relative to the respective baseline group. "All" includes all managers who mention either high prices or any alternative narrative. "Frequent" includes managers who mention either high prices or one of the relatively frequent (>5% frequency) alternative narratives. "Frequent without 'don't know'" is the same as the "Frequent" sample but excludes the "don't know" category. "Sales vol." includes managers who mention either high prices or narratives from the sales volume category. "Low prices" includes only managers who mention either high prices or the low prices narrative. All regressions control for manager characteristics (noncognitive skills, age, gender, experience), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, location type indicators) and district fixed effects. Corresponding regression coefficients are presented in Appendix Table B.1.

Robustness: Appendix B.1 presents similar results using alternative classification methods:

[24]A particular concern with the "don't know" narrative being in the omitted category is that this could be inversely related to cognitive skills due to measurement error from low survey effort—e.g., choosing "don't know" as the lowest-effort narrative response while also exerting low effort on the Raven test. As shown in Figure 2, however, the relationship of cognitive skills to the probability of mentioning high prices remains robust when excluding "don't know" responses (regression "Frequent without don't know").

Using a third RA to reconcile disagreements, rather than the authors; giving the rubric to a Large Language Model to do the classification rather than using RAs. Results also hold when restricting to managers who mention only one narrative, ruling out the possibility that high-skill managers mention high prices simply because they mention more causes overall (Figure B.4).

3.2 Self-reported pricing preferences and cognitive skills

We now examine whether cognitive skills also matter for self-reported pricing preferences. In our second survey wave, we asked managers a yes or no question about whether they "would like to charge lower prices than the default suggested by upper-level management" and if so, whether such price cuts would benefit "themselves," "the company," or "both." As mentioned in Section 2, the default suggestion of upper management is essentially always to price at the price ceiling. As shown in Figure 3, managers with lower cognitive skills are significantly more likely to prefer charging below the ceiling. This relationship remains significant after controlling for noncognitive skills, other manager characteristics, and station and market conditions (Appendix Figure B.5). Among managers favoring lower prices, 80% believe doing so "benefits the company," which fits our narrative evidence that lower-skill managers view price cuts as improving profitability.

Figure 3: Cognitive skills and preferences for lower prices

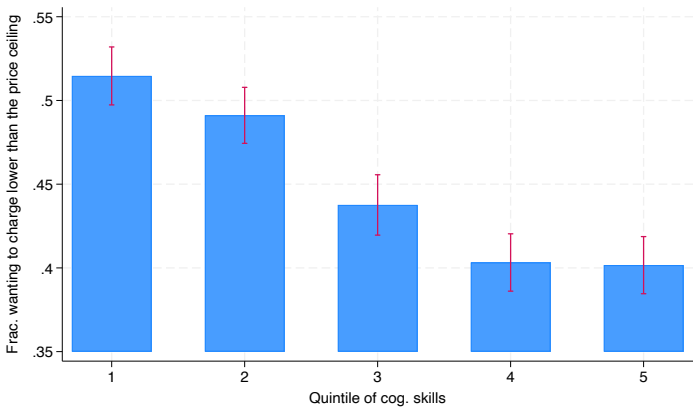

Notes: This figure reports the share of managers wanting to charge a lower fuel price than the default suggested by upper-level management, by quintile of cognitive skills (quintile 5 is the highest). Error bars indicate 95% confidence intervals.

3.3 Relationship of actual pricing to manager cognitive skills

Pricing decisions reflect a manager's need to respond to many factors in a changing environment (demand shocks, cost changes, competitor actions, and senior manager decisions), but we investigate whether cognitive skills also matter. We use monthly panel data on station-level pricing behavior. Because price ceilings vary and stations sell multiple fuel products

at different prices, the company computes a metric capturing overall monthly pricing: the sales-weighted average of each product's price relative to its ceiling. A ratio of 1 indicates pricing at the ceiling for all products. This price ratio incorporates all discounts, including reductions in posted prices, coupons, and promotions.

Figure 4: Pricing behavior as a function of cognitive skills

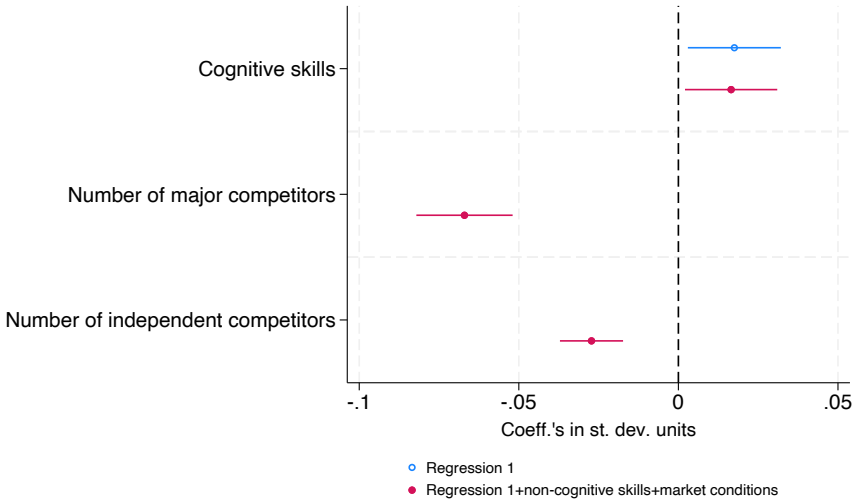

Notes: This figure reports coefficients from OLS regression on monthly price ratios, with 95% confidence intervals based on robust standard errors clustered at the station level. The dependent variable is the standardized monthly price ratio relative to the price ceiling. Regression 1 reports the coefficient on cognitive skills, controlling for manager demographics (age, gender, experience) and month-by-district fixed effects; Regression 2 additionally includes noncognitive skills and station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, location type indicators). See Appendix Table B.2 for underlying regression estimates.

Figure 4 shows that managers with lower cognitive skills set significantly lower prices relative to the ceiling, controlling for other manager traits, station and market characteristics, and time fixed effects. This finding aligns with our results on how low cognitive skills influence narratives about profit maximization and preferences for price cuts.

To put magnitudes in context, Figure 4 reports coefficients on our controls for number of independent and major competitors. A one-standard-deviation decrease in cognitive skills is associated with a price reduction roughly equivalent to three-fifths the effect of adding an independent competitor, or one-fourth the effect of adding a major competitor.[25] These magnitudes approximately double after accounting for attenuation from measurement error in cognitive skills—equivalent to adding more than one independent competitor or about half of a major competitor (see Section 2 for discussion of measurement error).

Robustness: As discussed in Section 2, all main results are robust to controlling for noncognitive skills individually rather than as a factor. As an example, Appendix Figure B.6 shows that

[25]Appendix Table B.2 shows a positive coefficient on the number of other stations from our partner company in the local market, consistent with high prices being easier to sustain between stations of the same company.

the coefficient on cognitive skills in regressions of the price ratio remains significant when controlling for noncognitive traits individually (individual elements are not significantly related to prices, with the exception of self-reported competitiveness, which is associated with significantly lower prices).[26]

3.4 Event study

In this section, we present results from an event study design comparing pricing at stations before and after a new manager arrives, and testing how cognitive skills of the new manager affect pricing. One purpose is to provide an additional check on causality. To be clear, we already have several reasons to believe our results are causal rather than driven by omitted variables: observable station and market characteristics are balanced by manager cognitive skills and we control for these characteristics in our regressions; we have evidence of plausible causal mechanisms in the form of different mental models of optimal pricing and different pricing preferences. In case there is measurement error in our controls for observables, however, or lack of balance in *unobservable* characteristics, we implement the event study to hold everything constant, comparing the same station and market environment before and after a new manager arrives. A second purpose of the event study is to examine the dynamics of how bounded rationality influences pricing over time: Whether the effect persists or diminishes with learning and feedback.

In our application, treated stations are those experiencing a manager change during the study period, with stations that never changed managers serving as controls. The identifying assumption is that, absent a manager change, treated and control stations would have followed parallel price trends. We compare treatment effects across stations receiving high- versus low-skill new managers, tracing how prices evolve differently depending on the new manager's cognitive skills.

This approach is robust to concerns about nonrandom assignment of managers to stations. If high-skill managers tend to be assigned to market conditions where charging high prices is optimal, such factors would be reflected in pre-period price levels, and these level differences are accounted for by comparing changes relative to the pre-period. Likewise, if assignment is nonrandom with respect to previous manager traits that affect pricing, this too would appear in pre-period prices and be addressed by the design. Offering additional assurance, we find no empirical evidence of selection when comparing previous and new manager cognitive skills: If upper management systematically assigned managers with certain cognitive skills to particular types of stations, or managers with certain cognitive skills self-selected into these types of stations, we would expect stations that previously had high-skill managers to also receive high-skill replacements; the correlation between new and pre-

[26]The measure asks, with an 11-point scale: "How competitive do you consider yourself to be? Please choose a value on the scale below, where 0 means not competitive at all and 10 means very competitive."

vious manager cognitive skills, however, is 0.019 ($p = 0.576$), indicating that assignment is essentially uncorrelated with cognitive skills.

We analyze the roughly 4,500 manager-change events in our dataset for which we know the traits of the new manager. While we sometimes know the traits of the previous manager, restricting to cases where both are known substantially reduces sample size, so we focus on how variation in new manager traits affects pricing after arrival. Since managers change stations at different times, we face a staggered treatment timing problem. Standard two-way fixed effects estimators can produce biased estimates in this setting because already-treated units serve as controls for newly-treated units. We therefore use the Callaway-Sant'Anna difference-in-differences (CSDID) estimator (Callaway and Sant'Anna, 2021), which computes group-time average treatment effects by comparing units treated at a particular time to never-treated units, then aggregates these effects.

Figure 5: Treatment effects of new managers on prices by cognitive skills

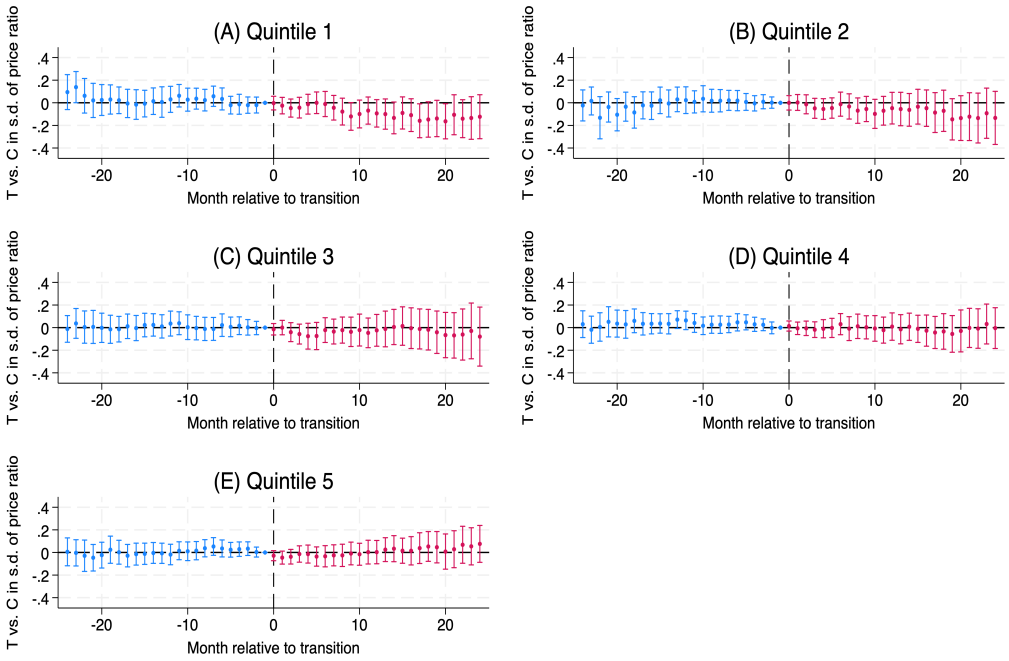

Notes: This figure reports treatment effects on price ratio in treated stations versus never-treated stations using the Callaway-Sant'Anna difference-in-differences (CSDID) estimator, by cognitive skills of the new manager. Treated stations are those that experienced a manager change; control stations are those without a manager change during the same period. Treatment effects are shown separately by quintile of cognitive skills (quintile 5 is the highest). Error bars indicate 95% confidence intervals. See Appendix Table B.3 for static average effects estimation.

Figure 5 presents the CSDID estimates by quintile of the new manager's cognitive skills. Parallel pre-trends hold across all five quintiles, validating our identification strategy. The results reveal a monotonic relationship: the lower the incoming manager's cognitive skills, the lower the treatment effect on prices. Averaging over the entire post-treatment period, both Quintile 4 and Quintile 5 show higher price levels than Quintile 1 (p-values of 0.071 and 0.056, respectively).

These differences become more pronounced and significant over time, with strikingly different time dynamics across skill levels. Quintile 1 managers charge progressively lower prices the longer they have been at a station. In contrast, Quintile 5 managers exhibit an initial price drop, perhaps reflecting experimentation as they learn about local market conditions such as competitor responsiveness, followed by price increases that eventually exceed those of the previous manager. These divergent patterns suggest that learning does not eliminate the impact of cognitive skills. Rather, both types experience similar environments but reach different conclusions, potentially because they interpret information through different mental models. For example, in the spirit of the level-k model in our conceptual framework, boundedly rational managers might interpret competitor price cuts as revealing lower costs, whereas sophisticated managers might interpret them as indicating boundedly rational competitors, making it optimal to keep prices high. In summary, the event study provides further evidence that cognitive skills cause persistent differences in pricing.

Robustness: Our results are robust to alternative estimators addressing the staggered diff-in-diff problem. Appendix Figures B.7–B.10 show consistent treatment effects across estimators proposed by De Chaisemartin and D'Haultfoeuille (2024), Sun and Abraham (2021), and Borusyak et al. (2024), as well as traditional two-way fixed effects.

4 Testing a potential mechanism: Mental models of competitor sophistication and price responses

This section investigates whether anticipating competitor sophistication and price responses is one mechanism linking cognitive skills to mental models of optimal pricing. To assess managers' mental models of competitors, we avoided direct questions about price cuts triggering competitor reactions, which could prompt respondents to consider such reactions even though they typically do not, or signal that surveyors expect them to do so. Instead, we designed three indirect approaches: (i) an abstract strategic game, in which choices reveal ability to anticipate competitor behavior; (ii) a survey question measuring confidence in ability to influence fuel sales, an indicator of believing in the efficacy of price cuts; (iii) a question measuring belief in the optimality of cutting price when being undercut. To validate whether cognitive skills, and our three measures of mental models of competitors, do capture anticipation of competitor price responses, we use a counterfactual prediction task in our third survey wave.

4.1 Money request game

As a measure of ability to anticipate competitors' strategic behavior, we implemented the money-request game.[27] The game was described as follows: "Suppose you are matched with another station manager to play a game. You and your opponent are each going to ask for an amount of money from a referee. The amount must be between $1 and $6. You will get the amount 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?"

The game is designed to have a salient, intuitive choice for players who are non-strategic: Choosing $6. The payoff-maximizing choice, however, depends on the empirical distribution of choices—if enough people choose $6 then $5 can be optimal, but if many choose $5, $4 may be optimal, and so on. Thus, making the optimal choice requires an accurate mental model of other managers' strategic reasoning. By design, the game has no pure-strategy Nash equilibrium, to help isolate ability to predict competitor behavior from ability to reason toward a pure-strategy NE.[28] While the money-request game differs from real-world competition—for example, it is one-shot rather than repeated—it provides a controlled measure of managers' ability to understand and predict the strategic behavior of other managers, holding constant all aspects of the strategic environment.

Figure 6 shows results from the money request game. Panel (A) presents the distribution of requests. The modal request is $5, consistent with many managers anticipating that others will act non-strategically and request $6. However, because so many managers request $5, the expected payoff is maximized by requesting $4 (requests below $4 yield lower expected payoffs). Thus, the modal manager makes a choice consistent with underestimating how many others choose $5. One might think requesting $4 is a matter of luck, but Panel (B) shows that average cognitive skills are significantly higher among managers choosing $4 compared to any other request, and Panel (C) shows the frequency of choosing $4 increases monotonically with cognitive skills. Regression analysis confirms that the probability of requesting $4 increases significantly with cognitive skills, controlling for other manager traits and market and station characteristics (first coefficient in Figure 7). To summarize, managers with higher cognitive skills are better able to anticipate competitor sophistication, consistent with our conceptual framework.

Robustness: As detailed in Appendix C.1, the money request game findings are robust across survey waves. In both the second and third waves, $5 was the modal request, $4 was the

[27] First proposed by Arad and Rubinstein (2012), this game has been a workhorse for laboratory research on bounded rationality and strategic competition. We used a simplified version developed by Fe et al. (2022). The original game allowed subjects to choose amounts between 11 and 20 Shekels, with a bonus of an additional 20 Shekels for requesting exactly 1 less than the opponent.

[28] Arad and Rubinstein (2012) argue this is an advantage over the classic beauty-contest game, which has a pure-strategy NE that empirically typically does not coincide with the payoff-maximizing choice. The mixed-strategy equilibrium involves probabilities of 0, 0, 0.4, 0.3, 0.2, and 0.1 for requests 1 to 6, respectively (Fe et al., 2022).

Figure 6: Behavior in the money request game and cognitive skills

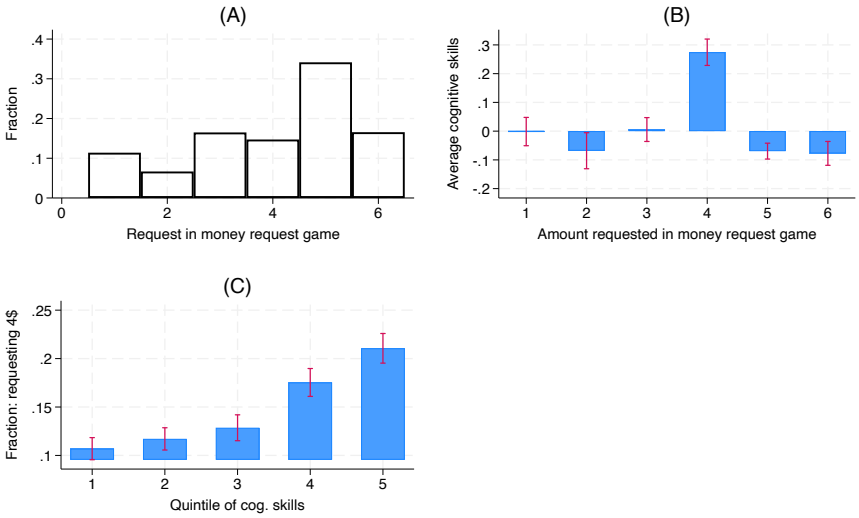

Notes: Panel (A) shows the distribution of requests in the money request game from the first survey wave. Panel (B) shows average cognitive skills for managers making each of the possible requests. Panel (C) shows the share of managers requesting $4, by quintile of cognitive skills (quintile 5 is the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

Figure 7: Mental models of competition and cog. skills: Controlling for other traits and market conditions

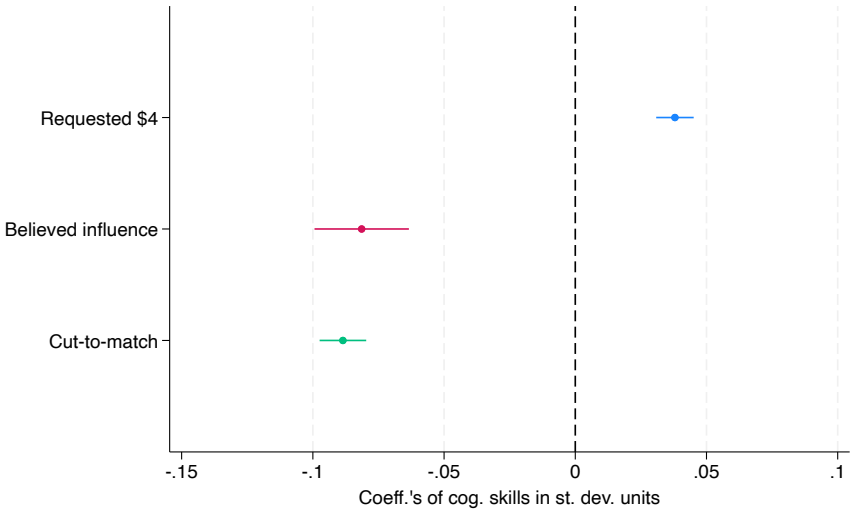

Notes: This figure reports coefficients from OLS regressions on three mental model measures, with 95% confidence intervals. The dependent variables are: requesting $4 in the money request game, believed influence over fuel sales, and believed optimality of the cut-to-match strategy. Coefficients shown are for cognitive skills. Each regression controls for manager characteristics (noncognitive skills, age, gender, experience), station and location characteristics (open 24 hours, whether the company rents the station, numbers and types of local competitors, location type indicators) and district fixed effects. See Appendix Table C.1 for underlying regression estimates.

payoff-maximizing choice, and those requesting $4 displayed the highest cognitive skills.[29]

[29]This adds to previous evidence that Raven test performance predicts choosing the payoff-maximizing request among children (Fe et al., 2022).

4.2 Beliefs about influencing fuel sales and optimality of cut-to-match strategy

We also designed two more concrete measures of managers' strategic thinking about their market environments, motivated by our initial in-person interviews (see Section 2). Some interviews revealed what we term an "attack" view: managers believed proactive price cuts could substantially increase sales, without recognizing that competitors might respond. Others revealed a "defense" view: Thinking it was optimal to cut price to match competitors' lower prices, without recognizing that doing so might trigger further competitor cuts.

We designed survey questions for each view. For attack we 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 of a station?" Based on our conceptual framework, we hypothesize that lower cognitive skills lead to *greater* confidence in the ability to influence fuel sales, because managers with lower skills do not fully anticipate how competitor price responses dampen the effects of their own price cuts. Although this inference—that confidence reflects not anticipating competitor price responses—is indirect, we validate this interpretation in the next subsection. We also asked about the ability to influence overall station performance and convenience store profits, as placebo tests, expecting that the effect of cognitive skills on anticipating competitor price reactions would matter mainly for confidence about fuel sales.

For the defense view, we asked a yes-or-no question: "Last month, you believed that it is optimal to cut price to match a competitor's price."[30] In our partner firm's terminology, cut-to-match refers to matching a competitor's lower price rather than maintaining the ceiling price. If the belief that cut-to-match is optimal reflects a failure to anticipate competitor responses, we would expect this belief to be more common among managers with lower cognitive skills.

Figure 8 shows how these measures relate to cognitive skills. Panel (A) shows that managers with lower cognitive skills are more confident in their ability to influence fuel sales, whereas cognitive skills are less strongly related to beliefs about influencing overall performance and largely unrelated to beliefs about influencing convenience store profits. This pattern is consistent with a mechanism operating specifically through the failure to anticipate competitor price responses in fuel markets. Panel (B) shows that the tendency to favor cut-to-match decreases monotonically across cognitive skill quintiles. Regression analysis (the second and third coefficients in Figure 7) confirms that these relationships are statistically

[30]The wording references the previous month because, to reduce the salience of our research interest in pricing, we embedded this in a management practices inventory based on McKenzie and Woodruff (2017), which asks about practices in a specific time frame. While our analysis focuses on cut-to-match, we also analyze two other price-related items: "Visited competitor's stations to see prices" and "Used a customized offer to attract new customers." Managers with high cognitive skills are less likely to report visiting competitor stations, consistent with placing less emphasis on price competition, and more likely to report attracting customers through customized offers rather than across-the-board price cuts.

Figure 8: Beliefs about manager influence and cut-to-match optimality, by cognitive skills

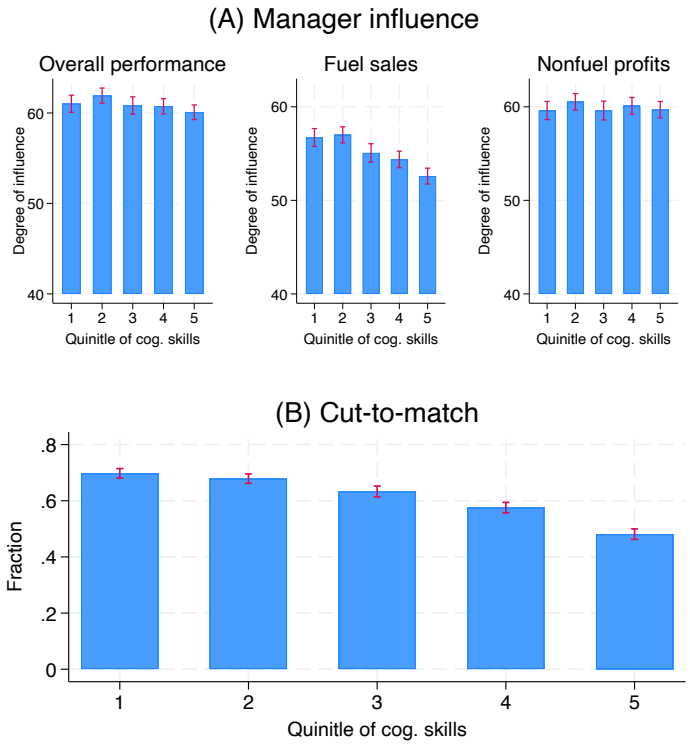

Notes: Panel (A) shows beliefs about manager influence measured on a scale from 0 (only external factors matter) to 100 (only manager matters). The three graphs show perceived manager influence on overall performance, fuel sales, and nonfuel profits, respectively, by quintile of cognitive skills (quintile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

significant and robust to controls for other traits as well as station and market characteristics. These findings are also in line with the assumption in the conceptual framework, that cognitive skills matter for mental models of competitors.

Robustness: In contrast to the abstract game, a concern with these more naturalistic measures is that results could reflect managers with different cognitive skills facing systematically different market conditions that shape learned beliefs, rather than bounded rationality generating divergent views of the same environment. However, as noted in Section 2, market conditions are comparable across cognitive skill levels, and as shown in Figure 7, results are robust to controls.

Because we test whether cognitive skills are related to three measures of mental models of competitors, a potential concern is multiple hypothesis testing. Applying the Bonferroni correction, however, all three coefficients in Figure 7 remain significant at $p < 0.001$.

4.3 Validation: Do cognitive skills and mental models of competitor behavior capture anticipation of price responses?

Our findings so far suggest that managers with lower cognitive skills are less likely to anticipate competitor price responses, but the evidence is relatively indirect. To gather more direct evidence, we first verify that competitors do respond to price cuts by our partner stations. We use data from one region with daily prices for both partner stations and all competitors. We define a price cut event as a partner station maintaining the same price for seven days and then reducing it by at least 10 cents. We employ a difference-in-differences approach comparing price responses of treated competitors to control competitors in the same district whose nearby partner stations maintained stable prices.

As shown in Figure C.3, competitors indeed respond to price cuts by lowering their own prices. Prior to price cuts, competitor prices show no differential trends, but immediately following partner station price cuts, treated competitors reduce their prices by approximately 5 cents on average—a reduction that persists over the following week. This establishes that competitor price responses are a salient feature of this market, which makes it consequential whether managers anticipate them. The figure also shows that competitors respond similarly regardless of the cognitive skills of the manager initiating the cut, meaning that managers across skill levels are exposed to comparable competitor price responses. However, managers may interpret competitor price changes differently depending on their cognitive skills, with low-skill managers viewing them as exogenous rather than as responses to their own pricing.

We next examine whether managers predict competitor price responses to price cuts, using a measure in our third survey wave implemented for a randomly selected sub-sample of roughly 6,000 managers.[31] 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 a competitor price cut in response to this earlier and particularly deep price cut.

We observe systematically different predictions of competitor prices depending on managers' cognitive skills. As shown in Figure 9, only 29% of managers in the bottom quintile of cognitive skills predict that Competitor 1 would respond to the hypothetical price cut by lowering its price. This share increases with cognitive skills, reaching 46% in the top quintile. A similar pattern holds for predictions of Competitor 2's price (Appendix Figure C.5).

[31] 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
May 11	5.01	2,118	4.28	3.98
May 12	5.01	1,981	4.28	3.98
May 13	5.01	855	4.28	3.98
May 14	5.01	1,530	4.28	3.98
May 15	5.01	3,334	4.28	3.98
May 16	4.81	1,600	4.28	3.98
May 17	4.81	2,841	4.28	3.98
May 18	4.81	1,528	4.29	3.98
May 19	4.81	1,764	4.29	3.98
May 20	4.81	4,279	4.29	3.98
May 21	4.81	679	4.08	3.98
May 22	4.81	2,014	4.08	3.98
May 23	5.01	2,179	4.08	3.98
May 24	4.81	2,203	4.08	3.98
	(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: The top panel reports actual historical data from a station of our partner company. The bottom panel presents the counterfactual scenario in which 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 real price and sales volume data at a gas station of your company over 14 days. During these 14 days, the price ceiling is constant at 5.01 per liter. 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 daily sales volume. The fourth and fifth columns show the prices of the first and second competitors, respectively."

Regression analysis confirms that the probability of predicting a competitor price reduction increases significantly with cognitive skills, controlling for other manager traits and station and market characteristics (Regression 1 in Figure C.6).

Examining the complete distribution of predicted prices provides additional insights (Appendix Figure C.4). Whereas most low-skill managers predict that Competitor 1's price would remain unchanged if our partner station lowered its price, those who do predict changes are roughly as likely to predict increases as decreases. This is consistent with beliefs that competitor prices move for reasons orthogonal to the station's own pricing decisions. In contrast, managers in the top cognitive quintile almost never predict price increases, have the lowest rate of predicting constant prices, and have the highest rate of predicting decreases. This implies a stark difference across cognitive skill levels in the tendency to predict that a price cut triggers competitor price cuts. Importantly, we do not claim that the predictions of high-skill managers are correct—indeed, the true counterfactual prices are

unobservable. Rather, our key finding is that managers with different cognitive skills make systematically different predictions after observing the same data for the same station, in a way that potentially explains differences in beliefs about optimal pricing and pricing decisions.

Figure 9: Cognitive Skills and Predicted Competitor Price Responses

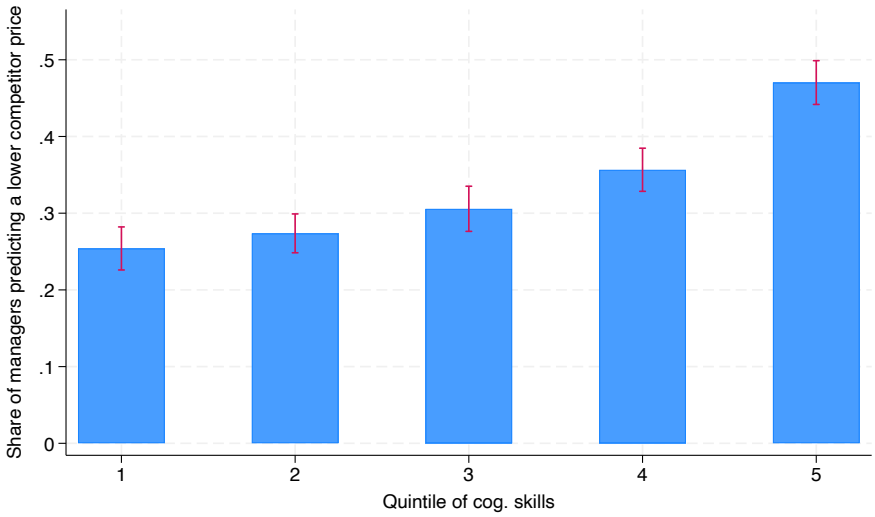

Notes: This figure reports the share of managers predicting a lower competitor price in response to a hypothetical price cut by a partner station, by quintile of cognitive skills (quintile 5 is the highest). Error bars indicate 95% confidence intervals.

We next regress an indicator for predicting a competitor price reduction on our three measures of mental models of competitors—the $4 money request, believed influence on fuel sales, and believed optimality of cut-to-match—controlling for manager and market characteristics (Appendix Figure C.6). All three are significantly associated with anticipating price cuts in the expected directions. These results validate our interpretation that the three measures capture how managers think about competitor price responses, and offer an explanation for why lower-skill managers tend to charge lower prices, in line with our framework—they are systematically less likely to think that their own price cuts will trigger competitor price cuts.

4.4 Mediation analysis

We next examine whether our three measures of mental models of competitors mediate the relationship between cognitive skills and pricing beliefs, preferences, and decisions. We begin with a standard mediation analysis based on the natural direct effect (Huber, 2019), adding the mental model measures as controls in regressions on cognitive skills. Appendix Figure C.7 shows that not anticipating competitor responses is associated with believing that low prices are optimal (Panel A), wanting to charge lower prices (Panel B), and actually charging lower prices (Panel C). Moreover, the coefficient on cognitive skills is substantially

attenuated when controlling for mental models.

Next, we examine whether mental models also mediate the causal effect of cognitive skills on prices in our event study. Appendix Figure C.8 presents event studies split by each mental model measure. For each measure, managers whose mental model implies that they anticipate competitor responses raise prices after arriving at a station, while those who do not anticipate competitor responses charge progressively lower prices. To provide additional evidence of mediation, we estimate the "controlled direct effect," which captures the effect of cognitive skills on prices that remains if we hold mental models of competitors fixed (Pearl, 2001; Acharya et al., 2016). When we stratify by mental models, the effect of cognitive skills is attenuated in all subsamples—those who anticipate competitor price responses and those who do not. Finally, we calculate the weighted controlled direct effect, which implies that our measures of mental models of competitors account for approximately 40% of the causal effect of cognitive skills on pricing.[32] Further details are provided in Appendix C.3.

Robustness: Our mediation analysis relies on model-free measures of the ability to anticipate competitor behavior, e.g., requesting \$4 in the strategic game. As a robustness check, we adopt a specific model of bounded-rationality, level-k, in line with the example model in our conceptual framework. Through this lens, choices in the money request game reveal a manager's level of reasoning. The game was designed to have a clear choice for non-strategic level-0 types, namely requesting \$6.[33] This implies that a request of \$5 reveals $k = 1$ and a request of \$4 reflects $k = 2$. Figure 6 shows that level-1 managers have lower cognitive skills than level-2, consistent with our model. Appendix C.4 provides further analysis, focusing only on the sample of individuals identified as level-1 versus level-2. Consistent with the model, level-2 types are significantly more likely than level-1 to predict that competitors will respond to a price cut by lowering their own prices (Figure C.16). Turning to actual pricing (Figure C.17), event studies show that level-1 types implement progressively lower prices over time, consistent with the prevalence of level-1 types as revealed by the game, combined with the price war dynamics predicted by the model when level-1 types interact. By contrast, level-2 types charge higher prices and maintain them at a high level, consistent with anticipating that price cuts trigger competitor price cuts.

[32]This approach estimates controlled direct effects at different values of the mediator, weighting them by the population distribution of mental models. The treatment effect gap between quintile 1 and quintile 5 managers is 0.068 standard deviations in the full sample. The weighted controlled direct effect is 0.041 standard deviations, implying mediation of 40%.

[33]Based on previous evidence, choices of \$3 and lower likely reflect random choice rather than higher levels of k. Even among university students—a relatively sophisticated subject pool—most individuals have $k < 3$, as revealed by open-text descriptions of reasoning; individuals choosing lower numbers tend to indicate they chose randomly (Arad and Rubinstein, 2012). More generally, the level-k literature has been summarized as finding that level-1 and level-2 are by far the most frequent types (Crawford and Iriberri, 2007).

5 Consequences for market prices, station profits, welfare, and measures of market power

5.1 Market prices and price wars

We first explore whether the lower price strategies of managers with low cognitive skills lead to lower market prices. We use data from one region with approximately 1,300 gas stations, including daily price information for all fuel products and competitor prices in each local market. We have already shown using these data that price cuts by our partner stations tend to trigger competitor price cuts (Appendix Figure C.3). As shown in Appendix D.1.1, our partner stations are also responsive to competitor-initiated price cuts, but with significant heterogeneity by cognitive skills: managers with below-median cognitive skills respond more strongly and immediately than those with above-median skills. As illustrated by the data shown in Table 3, cutting prices in response to being undercut can trigger further reductions by competitors.

Because low-skill managers are more likely to cut prices and more likely to respond aggressively to being undercut, it is plausible that they end up in more price wars. To analyze whether this occurs, we need to define a price war. There is no accepted quantitative definition in the literature, understandably so given that what constitutes a price war depends on the market context. A qualitative working definition proposed by Busse (2002) emphasizes mutual price cuts "significantly below the usually prevailing prices." We use the price ceiling as the prevailing price and define a price war as *mutual price cuts of at least 50 cents from the ceiling for 14 days or more,* where "mutual" means that both a partner station and at least one local competitor were involved. Appendix Figure D.2 illustrates a price war in our data. We observe 211 price wars across 1,324 non-monopoly fuel markets from 2018 to 2021. On average, price wars lasted 43 days (median 29 days), with partner stations lowering prices by 68 cents relative to the ceiling.

To analyze how price war involvement relates to cognitive skills, we aggregate to the manager level, using as the dependent variable the number of price wars in which the manager was involved. Managers in the lowest cognitive skill quintile experienced approximately 0.3 price wars over the three-year period; this frequency decreases steadily to less than 0.15 for those in the highest quintile—roughly twice as high for low-skill as for high-skill managers (Appendix Figure D.3). Regression analysis confirms that price war frequency is significantly negatively related to cognitive skills, controlling for other manager traits, station and market characteristics, and days of operation (Regression 1 in Appendix Figure D.4). In conclusion, while some price wars may be optimal—consistent with standard game-theoretic models of sustaining high prices—their substantially higher frequency among low-skill managers indicates that at least some are likely strategic mistakes.

Robustness: Results on price wars and cognitive skills are similar across three alternative specifications: lowering the threshold from 50 to 25 cents, requiring 21 rather than 14 days for sustained price cuts, and restricting to "isolated markets"—stations that share no competitors with other partner stations. The latter addresses the concern that price wars could be correlated across nearby markets when stations share a competitor (Appendix Figure D.4).

5.2 Profit consequences of bounded rationality and pricing

We now investigate how the pricing strategies of boundedly rational managers affect profits. Having established the impact of cognitive skills on pricing, we turn to the remaining link: How prices affect profits.

Estimating the causal effect of price on profit requires solving identification challenges, because the correlation between price and profit could reflect reverse causality: managers might adjust prices in response to profit realizations (although it is unclear why a low price should be a rational response to low profits). To address this, we estimate a structural demand model, using cost shifters as exogenous instruments for price. This allows us to estimate the demand function and compute how counterfactual price changes would affect producer surplus (PS), where changes in PS are equivalent to changes in profits.[34]

We model each station as facing a downward-sloping residual demand curve that incorporates both direct consumer responses to price changes and indirect effects through competitor reactions. Following a standard approach in the gasoline demand literature (Li et al., 2014; Coglianese et al., 2017), we assume constant elasticity of demand, specifying residual demand for station s as $q_s = A_s \cdot p_s^{-\varepsilon}$, where q_s is quantity sold, p_s is own price, $\varepsilon > 1$ is the price elasticity of demand, and A_s is a station-specific demand shifter. This elasticity captures the percentage change in quantity demanded in response to a one percent change in price, reflecting both consumers' reduced overall gasoline consumption and their substitution toward competing stations.

To identify the demand elasticity, we exploit exogenous variation in marginal costs driven by the government-mandated price ceilings. The ceilings are indexed to world oil prices, which is plausibly uncorrelated with local demand shocks. This instrument is particularly well-suited to our setting because our partner company is vertically integrated: under its internal accounting rules, the transfer price that determines each station's marginal cost moves one-for-one with the price ceiling. Crucially, the ceiling is adjusted every ten working days, providing high-frequency variation that allows us to include flexible day-of-week, month, and year fixed effects while retaining sufficient variation for identification. Specifically, for each fuel product j at station s, we log-linearize the demand equation and instru-

[34]The discrepancy between profits and PS—namely, fixed costs—is differenced out when calculating the change in profits.

ment for the log of the product-level price p_{sjt} using the log of the corresponding ceiling $Ceiling_{jt}$, controlling for station-product fixed effects and the time fixed effects described above. Pooling across all station-product observations, the 2SLS estimate yields $\hat{\varepsilon} = 1.65$. Further details are provided in Appendix D.2.

With an estimate of the elasticity $\hat{\varepsilon}$, we can calculate the profit implications of cognitive skill differences. We focus on an average station in our sample, using mean price $\bar{p}$ and sales volume $\bar{q}$ as the baseline. Given these values, we can back out the demand shifter as $\bar{A} = \bar{q} \cdot \bar{p}^{\hat{\varepsilon}}$. The remaining input needed for calculating profit losses is average marginal cost. A key advantage of our setting is that we directly observe this object: because our partner company is vertically integrated, the internal transfer price provides us with rare access to administrative data on marginal costs. This contrasts with the standard empirical industrial organization approach, which infers marginal costs from pricing behavior under the assumption that firms price optimally (Berry et al., 1995; Nevo, 2001). This assumption becomes problematic in the case of suboptimal pricing of boundedly-rational managers. With these inputs and the measured relationship between cognitive skills and pricing from Section 3, we can compute the counterfactual price, sales, and profits that would result from replacing the current manager with one who has lower cognitive skills. We find that a one-standard-deviation decrease in cognitive skills leads to price cuts that reduce daily profits by 2.6%. This indicates that lower price strategy of managers with low cognitive skills is a mistake from the perspective of the firm's profits.

Robustness: Our structural estimation assumes constant elasticity of demand, which is standard in the gasoline demand literature and provides a tractable framework for estimating the residual demand curve. In Appendix D.3, we show, however, that our results are robust to an alternative specification using linear demand. This alternative allows us to separately identify own-price and cross-price effects. Under this specification, the estimate of demand elasticity for the average station is -1.63, very similar to our result with constant elasticity demand and thus implying a comparable 2.6% drop in daily profits.

5.3 Long-run benefits of lower prices?

A potential limitation of our structural estimation is that, while it can assess short-run effects of price cuts on profits, it does not account for possible long-run benefits such as attracting loyal customers or deterring potential entrants.

To assess whether such benefits might exist, we first turn to survey evidence from district managers. As shown in Table 4, district managers generally believe that station managers tend to cut prices if given more autonomy. When then asked whether prices chosen by station managers would be too high, about right, or too low, roughly 75% answered too low. When asked in an open-ended question whether station managers should be "allowed to actively participate in price competition and charge lower prices than competitors," 74%

indicated such strategies are not beneficial. Of these, 71% took a stronger position, stating that aggressive price competition is always a bad idea. Of those offering specific explanations, half stated that lower prices directly reduce profits, while the other half cited the risk of triggering price wars. These findings indicate that most district managers believe that low prices reduce profits and do not consider low-price strategies profitable in the long run.

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: This table reports results from the survey of district-level managers.

As additional evidence, we examine the relationship between cognitive skills and realized profits. If low-price strategies yielded long-run benefits, these should eventually materialize as higher profits. Appendix Figure D.5 shows the opposite: managers with higher cognitive skills earn significantly higher total profits for both fuel and nonfuel categories. This indicates that any long-run benefits from low-skill managers' pricing strategies are insufficient to offset the short-run losses documented in our structural analysis.

5.4 Welfare consequences beyond producer surplus

We now examine how cognitive skill differences affect consumer surplus (CS) and dead weight loss (DWL). Further details are provided in Appendix D.5. The change in CS is measured by the area to the left of the demand curve between the two price levels implied by cognitive skill differences. Given downward-sloping demand and price above marginal cost, a reduction in price necessarily increases CS and reduces DWL. Our estimates, however, allow us to calculate the magnitude. We calculate an increase in CS of 1.3% from a one-standard-deviation decrease in manager cognitive skills. Notably, CS is roughly five times as large as PS in levels, so a 1.3% increase in CS more than offsets the decrease in PS and implies an increase in total surplus. Correspondingly, the lower price implies a substantial 6.2% reduction in DWL.

Reducing bounded rationality with artificial intelligence The preceding analysis can also be read in reverse: what are the implications of *increasing* managerial cognitive skills? This perspective is relevant for understanding the welfare effects of AI-based algorithmic pricing systems. If such systems raised the effective level of cognitive skills by one standard deviation at an average station, PS would increase by approximately 3%, while CS would decrease by 1.3% and DWL would increase by 6.2%. Given that our partner company and one other

major firm together control approximately two-thirds of the market, widespread adoption of algorithmic pricing could meaningfully reduce market efficiency, with producers capturing additional surplus at consumers' expense.

5.5 Implications for measuring market power

In standard approaches to measuring market power, researchers either observe marginal cost or back it out by assuming firms price optimally. Our findings suggest that cognitive skills introduce a third source of price variation that could be mis-attributed under either approach. Consider first the case where marginal cost is observed, as in our setting due to vertical integration. A researcher observing price variation across stations with the same marginal cost might attribute this to differences in market power or demand elasticity. However, our results show that cognitive skills alone can generate substantial price variation: a one-standard-deviation reduction in cognitive skills reduces the markup $\frac{p-MC}{p}$ by approximately 3 percent (0.29 percentage points). Thus, what appears to be variation in market power may partly reflect variation in managerial cognitive skills. The more common case in empirical industrial organization is that marginal cost is *not* observed. Researchers typically assume optimal pricing and back out marginal cost from the first-order condition. If a researcher observes two stations with similar market conditions but different prices, the standard approach attributes this entirely to differences in marginal cost. Our results suggest this attribution may be incorrect: the price difference could reflect differences in cognitive skills, with both stations facing the same marginal cost but the lower-priced station failing to fully exploit its market power.

6 Conclusion

This paper provides the first comprehensive field evidence on whether and how bounded rationality among firm managers affects strategies in repeated price competition. Using data from a company with over 20,000 gas station managers, including detailed measures of cognitive skills, mental models, and pricing decisions, we document three sets of findings. First, managers with lower cognitive skills exhibit a systematic tendency toward lower prices in their mental models of optimal pricing and charge lower fuel prices—an effect similar in magnitude to adding three-fifths of a competitor, which persists over time despite high stakes and repeated feedback. Second, one underlying mechanism is failure to anticipate competitor sophistication and price responses. Third, the pricing decisions of boundedly rational managers have substantial economic consequences: triggering competitor price cuts, increasing price war frequency, reducing producer surplus, increasing consumer surplus and market efficiency, while also biasing standard measures of market power.

These findings have important implications for theory and policy. They demonstrate that high stakes and learning opportunities may not eliminate the impact of cognitive constraints on pricing in repeated price competition. Because these constraints can lead to systematic effects rather than just noise, firms seeking to improve pricing may need to reshape managers' mental models, and switching to sophisticated algorithmic pricing could substantially affect market prices and efficiency. By limiting anticipation of competitor price responses, bounded rationality may be a force for lower prices in any market characterized by repeated price competition and market power, and phenomena like price cuts and price wars may at least sometimes reflect mistakes rather than optimal strategies.

One further implication concerns the emergence of price coordination in markets. While coordination on high prices has been documented in a range of markets, our findings may help explain why such coordination does not always emerge and, when it does, can take substantial time.[35] The price ceiling in our setting arguably makes coordinating on high prices easier, suggesting our results may provide a lower bound on the effect of bounded rationality on prices and efficiency.

Our results also speak to a tradeoff not previously emphasized in the literature: the implications of bounded rationality for organizational structure and delegation. While managers in our partner company have substantial pricing autonomy, retail gas companies in some markets centralize pricing, and implement uniform pricing across locations (Foros and Steen, 2013). More broadly, DellaVigna and Gentzkow (2019) show that uniform pricing across locations is common in US retail chains and can be very costly due to forgone price discrimination. A potential explanation for why firms nonetheless adopt uniform pricing is that decentralized pricing is even more costly when managers price suboptimally. That is, retail chains may accept the cost of uniform pricing to avoid the larger cost of boundedly rational local pricing. Thus, our findings are relevant even for markets with centralized pricing—such structures may themselves be response to bounded rationality.

While our primary focus is on cognitive skills, 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 traits other than cognitive skills. Noncognitive skills and experience predict higher confidence in ability to influence fuel sales and in the optimality of price matching—that is, the same direction as *low* cognitive skills (Table C.1). However, noncognitive skills and experience predict a greater likelihood of mentioning high prices in profit narratives (Table B.1), similar to *high* cognitive skills. Since these traits do not lead to more accurate mental models of competitor behavior, it seems they are related to mentioning high prices for different reasons than high cognitive skills, perhaps recognizing that the ability to keep prices high can be beneficial without successfully doing so in practice. This interpretation is supported by the finding that these traits are unrelated to, or even neg-

[35]For example, Byrne and de Roos (2019) show that coordination on high prices takes about three years to emerge in an urban retail fuel market.

atively correlated with, actual fuel prices (Table B.2). Overall, the relationships between these traits and mental models are not as robust and consistent as those for cognitive skills, and noncognitive skills and experience are not substitutes for cognitive skills in pricing.

Our findings raise questions for future research. What interventions might improve pricing decisions by boundedly rational managers, given that differences in mental models persist despite access to data and experience? Does social learning within firms reinforce incorrect mental models, or do more sophisticated managers' views prevail? Finally, what other mechanisms link cognitive skills to pricing? We find evidence for underappreciation of competitor responses, but cognitive skills may also influence understanding of consumer responses. In follow-up work, we plan to explore whether lower-skill managers overestimate consumer responsiveness to price cuts—for example, by failing to recognize that some sales responses reflect intertemporal substitution rather than new customer acquisition.

References

Acharya, Avidit, Matthew Blackwell, and Maya Sen, "Explaining causal findings without bias: Detecting and assessing direct effects," *American Political Science Review*, 2016, *110* (3), 512–529.

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.

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

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.

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, "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," *Journal of Political Economy*, 2024, *132* (3), 723–771.

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, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

_ , 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.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

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.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of econometrics*, 2021, *225* (2), 200–230.

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.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Difference-in-differences estimators of intertemporal treatment effects," *Review of Economics and Statistics*, 2024, pp. 1–45.

Coglianese, John, Lucas W Davis, Lutz Kilian, and James H Stock, "Anticipation, tax avoidance, and the price elasticity of gasoline demand," *Journal of Applied Econometrics*, 2017, *32* (1), 1–15.

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.

Crawford, Vincent P and Nagore Iriberri, "Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Danz, David N, Dietmar Fehr, and Dorothea Kübler, "Information and beliefs in a repeated normal-form game," *Experimental Economics*, 2012, *15* (4), 622–640.

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

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.

Duffy, John and Rosemarie Nagel, "On the robustness of behaviour in experimental 'beauty contest'games," *The Economic Journal*, 1997, *107* (445), 1684–1700.

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*.

Foros, Øystein and Frode Steen, "Vertical Control and Price Cycles in Gasoline Retailing," *The Scandinavian Journal of Economics*, 2013, *115* (3), 640–661.

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.

_ **and Yaroslav Rosokha**, "Beliefs, learning, and personality in the indefinitely repeated prisoner's dilemma," *American Economic Journal: Microeconomics*, 2024, *16* (3), 259–283.

_ **and** _ , "Identifying level-k reasoning in repeated games: Strategies, beliefs, and cognitive ability," *Working paper*, 2026.

_ **, Victoria L Prowse, and J Lucas Reddinger**, "Dynamic investment in teamwork skill: Theory and experimental evidence," *Available at SSRN 5433715*, 2025.

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 Botao Yang, "Are all managers created equal?," *Journal of Marketing research*, 2009, *46* (5), 612–622.

_ **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," *Quantitative Marketing and Economics*, 2024, *22* (3), 223–255.

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

Hansen, Karsten T., Kanishka Misra, and Mallesh M. Pai, "Frontiers: Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms," *Marketing Science*, 2021, *40* (1), 1–12.

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, Fernando Luco, Steven L Puller, and Dongni Zhu, "Does strategic ability affect efficiency? Evidence from electricity markets," *American Economic Review*, 2019, *109* (12), 4302–4342.

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

Huber, Martin, "A review of causal mediation analysis for assessing direct and indirect treatment effects," 2019.

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.

Kaplan, Robert S and David P Norton, "The Balanced Scorecard: Measures that Drive Performance," *Harvard Business Review*, 1992, *70* (1), 71–79.

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.

Li, Shanjun, Joshua Linn, and Erich Muehlegger, "Gasoline Taxes and Consumer Behavior," *American Economic Journal: Economic Policy*, 2014, *6* (4), 302–342.

List, John A, Ian Muir, Devin Pope, and Gregory Sun, "Left-digit bias at lyft," *Review of Economic Studies*, 2023, *90* (6), 3186–3237.

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

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.

Nevo, Aviv, "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 2001, *69* (2), 307–342.

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

Olea, José Luis Montiel, Pietro Ortoleva, Mallesh M Pai, and Andrea Prat, "Competing models," *The Quarterly Journal of Economics*, 2022, *137* (4), 2419–2457.

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

Pearl, Judea, "Direct and indirect effects," *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, 2001, pp. 411–420.

Proto, Eugenio, Aldo Rustichini, and Andis Sofianos, "Intelligence, personality, and gains from cooperation in repeated interactions," *Journal of Political Economy*, 2019, *127* (3), 1351–1390.

_ , _ , and _ , "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, "Conjectures, firm characteristics, and market structure: An empirical assessment," *International Journal of Industrial Organization*, 1986, *4* (4), 347–369.

_ , "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.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of econometrics*, 2021, *225* (2), 175–199.

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.