

Can Better Information Reduce Efficiency? Information Synergies in Financial Mergers*

Marco Cosconati¹, Fan Wu², and Yi Xin³

¹IVASS, Bank of Italy

²Peking University, HSBC Business School

³California Institute of Technology

October 21, 2026

Abstract

Mergers among financial institutions can generate information synergies beyond conventional cost synergies. We develop a framework to study how such information synergies affect competition and market outcomes, and show that improved screening by merging firms can distort the allocation of consumers across firms, thereby reducing allocative efficiency. We quantify these effects through merger simulations for the Italian auto insurance market. We find that mergers generate substantial informational advantages, allowing merged firms to cream-skim low-risk consumers and reduce their average expected claim payouts by 7.2–15.5%. At the market level, however, these private gains can come at an efficiency cost: for some mergers, the resulting reallocation of consumers increases the average cost of serving the market by up to 1.4%, or approximately 12 euros per contract.

Keywords: Financial Intermediation, Mergers and Acquisitions, Information Synergies, Insurance Competition, Merger Simulation, Distributional Effects, Digital Economy, Big Data and AI Analytics.

*Cosconati: IVASS and Bank of Italy, via del Quirinale, 21, 00187, Roma. Email: marco.cosconati@gmail.com. Wu: Peking University HSBC Business School, University Town of Shenzhen, China. Email: fanwu@phbs.pku.edu.cn. Xin: Division of the Humanities and Social Sciences, California Institute of Technology, 1200 East California Blvd, MC 228-77, Pasadena, CA 91125. Email: yixin@caltech.edu. We thank seminar and conference participants at Stony Brook University, University of Calgary, University of Hong Kong, and 2026 Asia Meeting of the Econometric Society. The opinions and errors in this paper are solely those of the authors and do not reflect the views of IVASS and the Bank of Italy.

JEL Codes: D82, G22, G28, L13, L41.

1 Introduction

Information is a central input into financial intermediation. Banks, insurers, and other financial institutions invest heavily in data and analytics to assess risk, screen customers, set prices and contract terms, and determine which market segments to serve. Advances in big data, machine learning, and artificial intelligence have increased both the value of proprietary information and the returns to combining information from different sources. As a result, mergers among financial institutions can create an important form of information synergy: by integrating complementary datasets and analytical capabilities, the combined firm may become substantially better informed about its customers than its competitors.

Does such an improvement in information necessarily make financial markets more efficient? At first glance, the answer might appear to be yes. In markets where asymmetric information plays a central role, better information allows financial institutions to assess customer risk more accurately. This logic underlies much of the interest in technologies that improve screening and in policies that facilitate information sharing, such as credit registries and open banking. Yet a firm's informational advantage can also reshape which customers it chooses to compete for. By changing *who serves whom* in the market, a financial merger can have ambiguous effects on the efficiency of market allocation.

This paper studies the equilibrium and welfare consequences of information synergies generated by mergers among financial institutions. Our central insight is that an improvement in a firm's screening technology following a merger can distort the allocation of consumers across firms, thereby reducing allocative efficiency. A merger can enable the combined firm to identify particularly attractive low-risk consumers and compete aggressively for them. But the merged firm need not be the lowest-cost provider for those consumers. If its cost advantage instead lies in serving higher-risk customers, its new informational advantage can create a misalignment between firms and consumers that increases total market costs. This mechanism distinguishes information synergies from the cost synergies traditionally associated with mergers, which directly lower production costs, for instance through reduced transportation costs or

improved managerial efficiency.¹

Among financial markets, insurance provides a particularly natural setting for studying these issues. The insurance industry is inherently data-intensive: insurers’ core activity consists of collecting, processing, and analyzing information to assess risks and price contracts accordingly. This process relies on extensive data on consumer characteristics and detailed historical claims records, which are increasingly integrated into advanced AI-based underwriting models. Moreover, the industry has recently experienced significant merger activity, often justified by the integration of complementary data and analytic capabilities.² Understanding how these merger-driven information synergies affect competition and consumer welfare has therefore become increasingly important for policy.

In this paper, we begin by developing a stylized model to analyze the key mechanisms through which information synergies generated by mergers affect market efficiency and consumer welfare in insurance markets. Specifically, we consider a setting with three risk-neutral firms competing in prices for consumers. Consumers are either high- or low-risk types. Firms receive imperfect signals about consumers’ risk types and may differ in their costs of serving high- and low-risk consumers. A merger allows the combined firm to integrate the signals of both parties, thereby improving its information precision.³ In this simplified model, we abstract from market power and product differentiation, assuming that firms price at expected cost and consumers choose the lowest-priced contract. This setup allows us to isolate the welfare consequences of information-driven selection from those arising from strategic pricing distortions and demand-side frictions.

The model delivers two main results. First, information synergies can either in-

¹For example, [Ashenfelter et al. \(2015\)](#) find that the 2008 Miller–Coors merger reduced shipping costs, while [Braguinsky et al. \(2015\)](#) show that acquisitions in the Japanese cotton spinning industry improved inventory management, accounts receivable, and capacity utilization.

²For example, Progressive’s acquisition of ASI was motivated in part by the desire to integrate auto and homeowners data in order to develop a unified household-level risk and pricing framework. Similarly, a proposed 2021 merger between two major insurance brokers, Aon and Willis Towers Watson, ultimately blocked by the DOJ, aimed to combine proprietary datasets and modeling techniques to strengthen data-driven analytics. In the Italian auto insurance market that we study, AXA recently acquired Prima, a technology-driven insurer known for its proprietary pricing algorithms.

³Our paper focuses on horizontal mergers between insurance companies. By contrast, other recent work examines mergers across data-connected markets (e.g., insurance and wearable health devices), in which the merged firm leverages data collected in one market to enhance product customization and price discrimination in another ([Chen et al., 2022](#); [Herresthal et al., 2023](#); [De Cornière and Taylor, 2024](#)).

crease or decrease allocative efficiency. Following a merger, the combined firm can identify particularly low-risk consumers more precisely and offer them more favorable prices. This informational advantage allows the merged firm to cream-skim the safest consumers. When the merged firm also has a comparative cost advantage in serving these consumers, the resulting reallocation improves efficiency. But when its cost advantage lies instead in serving high-risk consumers, cream-skimming moves low-risk consumers away from the more efficient provider, thereby increasing total market costs. Better information at the firm level can therefore reduce efficiency at the market level.

Second, the welfare consequences of information synergies are heterogeneous across consumer risk types and can be counterintuitive. Consumers who are identified as particularly low risk benefit from more accurate risk classification. In contrast, consumers in the intermediate-risk segment are worse off because they are separated from the low-risk group and may instead be pooled with higher-risk consumers. Such pooling arises when rival firms have cost advantages that enable them to compete for intermediate- and high-risk consumers. When this occurs, high-risk consumers are effectively cross-subsidized and face lower premiums relative to the pre-merger equilibrium, leading to an increase in their consumer surplus. This result stands in contrast to the standard “information effect” in [He et al. \(2023\)](#), whereby improved information benefits low-risk consumers but harms high-risk consumers.

To assess the quantitative importance of the mechanisms highlighted by the stylized model, we turn to a richer empirical model of competition in the Italian automobile insurance market. The framework builds on [Cosconati et al. \(2025\)](#), which studies insurer competition and estimates firms’ heterogeneous risk-rating precision, cost structures, and consumer demand in this market.⁴ The empirical model extends several key elements of the stylized model to better capture relevant institutional and market features. In particular, it allows for imperfect competition, rich heterogeneity in firms’ information precision and cost structures, a continuous distribution of consumer risk, and a more realistic demand system with differentiated products.

Using estimates from the Italian auto insurance market, we simulate all 45 potential pairwise mergers among the 10 largest insurers and solve for the resulting

⁴Relatedly, [Crawford et al. \(2018\)](#) investigate imperfect competition among banks in Italian loan markets, and [Blickle et al. \(2025\)](#) study credit market competition when specialized lenders have an informational advantage regarding borrowers in their focal industries.

market equilibria, allowing prices, consumer choices, and hence the allocation of risks across insurers, to adjust. Importantly, these simulations account for both changes in market structure and the information synergies generated by each merger. We find that information synergies indeed enable the merged firm to cream-skim low-risk consumers, reducing its average expected claim payouts by 7.2 to 15.5 percent relative to the pre-merger benchmark. Moreover, our simulations indicate that these informational advantages can generate significant efficiency losses. Among the 45 mergers, 20 percent lead to an increase in the total cost of serving the market, by up to 1.4 percent, or approximately 12 euros per contract.⁵ Notably, we find that these efficiency losses arise when the merged firm is particularly efficient at processing claims but are not necessarily associated with larger increases in market concentration, as measured by the HHI. This finding suggests that standard concentration measures may fail to capture important efficiency considerations in merger analysis. It underscores the need for competition authorities to assess changes in firms’ informational capabilities in financial mergers alongside conventional concerns about market concentration.

In terms of consumer welfare, our merger simulations show that high-risk consumers disproportionately benefit from information synergies, while the effect on low-risk consumers is modest. Across the 45 merger cases we consider, high-risk consumers experience an average increase in consumer surplus of 4.75 percent, whereas low-risk consumers experience an average decrease of 0.16 percent. Information synergies have consequences beyond aggregate efficiency and average prices—the traditional focus of antitrust analysis: they also reshape the distribution of surplus across consumers with different risk profiles. Determining how to balance these trade-offs across consumer groups ultimately depends on the objectives of merger policy, including the relative emphasis placed on efficiency, equity, and consumer protection.

Related Literature Our paper is related to several strands of literature. First, we contribute to the literature on mergers among financial institutions. Existing work has studied how consolidation affects market power, operating efficiency, credit supply, diversification, and financial stability (see, e.g., the survey articles by [Berger et al. \(1999\)](#) and [DeYoung et al. \(2009\)](#)). A smaller literature considers the informational

⁵This effect is economically meaningful: it is comparable in magnitude to the efficiency gains from establishing a centralized risk bureau that equalizes access to information across competing insurers, as estimated by [Cosconati et al. \(2025\)](#).

consequences of bank mergers. [Hauswald and Marquez \(2006\)](#), for example, develop a theoretical model of competition in credit markets and shows that banks may have incentives to merge when consolidation enables the coordination of screening activities across branches. Using data from the Italian loan market, [Panetta et al. \(2009\)](#) provide empirical evidence that bank mergers improve screening ability.

Our paper contributes to this literature by analyzing how merger-driven information synergies affect competition and consumer welfare in financial markets and by quantifying the magnitude of these effects using a structural model of the Italian automobile insurance market. We show that information synergies can generate economically significant losses in allocative efficiency and redistribute surplus across consumer groups, potentially benefiting higher-risk consumers.⁶

Second, our paper is related to the literature on information sharing in oligopoly. The classic theoretical literature studies firms’ incentives to exchange information about market demand or costs (e.g., [Gal-Or, 1985](#); [Shapiro, 1986](#); [Kühn and Vives, 1995](#); [Banal-Estañol, 2007](#)), while empirical work documents information sharing through trade associations, public earnings announcements, and other disclosure mechanisms (e.g., [Christensen and Caves, 1997](#); [Doyle and Snyder, 1999](#); [Aryal et al., 2022](#)).⁷ More closely related to our setting in financial markets, [Pagano and Jappelli \(1993\)](#) study the sharing of borrower information among banks and show that centralized credit bureaus can improve market performance in the presence of moral hazard and adverse selection. [He et al. \(2023\)](#) study open banking in a model of competition between a bank and a fintech, where data sharing improves the fintech’s screening ability. They show that competition intensifies as lenders’ screening abilities become more similar and weakens as they diverge.

Our setting differs in an important way. Information-sharing arrangements such as credit registries or open banking can reduce informational asymmetries across competing firms, whereas a merger can *widen* them.⁸ We show that such synergies

⁶This finding echoes the results in [Allen et al. \(2014\)](#), who show that mergers in the mortgage market disproportionately affect consumers with lower search costs.

⁷Empirical studies also document firms sharing pricing information through trade journals ([Marshall et al., 2008](#)), private information platforms ([Borenstein et al., 2004](#); [Miller, 2010](#); [Byrne et al., 2025](#)), and price transparency or disclosure policies ([Rossi and Chintagunta, 2016](#); [Luco, 2019](#); [Ater and Rigbi, 2023](#)).

⁸[He et al. \(2023\)](#) also analyze a scenario in which open banking “over-empowers” the fintech, thereby widening the gap in screening ability between the two lenders. [Byrne et al. \(2025\)](#) study an antitrust case in which one firm loses access to high-frequency, granular rival pricing data, resulting in an asymmetric information sharing among competitors.

can either increase or reduce market efficiency, depending on how the merged firm’s information technology interacts with its cost structure.

Third, we contribute to the broader merger literature by formalizing information integration as a distinct source of merger synergies. Since the seminal work of [Williamson \(1968\)](#) and [Farrell and Shapiro \(1990\)](#), the empirical merger literature has primarily focused on the trade-off between cost efficiencies and concentration effects.⁹ We show that information synergies have fundamentally different welfare implications from conventional cost synergies and therefore warrant distinct consideration in merger analysis, especially given their growing relevance in financial markets and other data-intensive industries.

The remainder of the paper is organized as follows. Section 2 presents a stylized model that illustrates the key mechanisms through which information synergies arising from mergers affect market efficiency and consumer welfare. Section 3 describes the institutional setting of the Italian insurance market and introduces the empirical framework. Section 4 presents the merger simulation results and Section 5 concludes. The proofs of the theoretical results are provided in Appendix A.

2 A Stylized Model

To highlight the key mechanisms through which information synergies generated by mergers can influence market efficiency and consumer welfare, we consider the following stylized model. We assume two types of consumers, $\theta \in \{l, h\}$, where l and h denote low- and high-risk types, respectively. Let μ_0 denote the share of high-risk consumers in the market, which is common knowledge.

In this market, there are three risk-neutral firms, $j \in \{1, 2, 3\}$. We assume that insurance is mandatory, so consumers must purchase a product from one of the three firms. This assumption is motivated by the empirical context we examine later in Section 3. Firms 1 and 2 are symmetric in that they share the same information technology and cost structure. Each firm $j \in \{1, 2\}$ can acquire a private signal $s_j \in \{l, h\}$ about a consumer’s type. We assume that s_1 and s_2 are conditionally

⁹A large body of empirical research documents cost reductions and managerial improvements following mergers (e.g., [Focarelli and Panetta, 2003](#); [Jeziorski, 2014](#); [Ashenfelter et al., 2015](#); [Braguinsky et al., 2015](#); [Grieco et al., 2018](#); [An and Zhao, 2019](#); [Eliaison et al., 2020](#); [Elliott et al., 2025](#)). Also see the review articles by [Röller et al. \(2006\)](#) and [Asker and Nocke \(2021\)](#) for comprehensive discussions of efficiency and concentration effects in merger analysis.

independent given the true type and coincide with the true type with probability q . The parameter q captures the information precision of Firms 1 and 2. In contrast, Firm 3 is assumed to be fully uninformative about the consumer's type. Regarding costs, Firm j 's cost of serving a type- l and a type- h consumer is given by c_j^l and c_j^h , respectively. High-risk consumers are more costly to serve, so that $c_j^l < c_j^h$. Under the symmetry assumption for Firms 1 and 2, we have $c_1^l = c_2^l$ and $c_1^h = c_2^h$.

Despite the simplicity of the stylized model, characterizing the full equilibrium is nontrivial, as firms' pricing strategies depend on posterior beliefs, which in turn depend on competitors' strategies and endogenous consumer choices. To focus on the key mechanisms, we impose the following simplifying assumptions, all of which will be relaxed in the empirical framework in Section 3.

Importantly, we assume that firms have no market power; that is, they set prices equal to expected costs, conditional on their available information and on being selected by the consumer. The pricing strategy of Firm $j \in \{1, 2\}$ is therefore

$$p_j(s_j) = \mathbb{E}(c_j | s_j, D = j),$$

where D denotes the consumer's contract choice. Accounting for consumer choice captures the winner's curse effect (i.e., the possibility that winning a consumer reflects the rival firm's receipt of an unfavorable signal). Since Firm 3 does not receive any signals, it sets its price equal to expected cost conditional on being chosen by the consumer, i.e.,

$$p_3 = \mathbb{E}(c_3 | D = 3).$$

The no-market-power assumption provides a natural benchmark for welfare analysis of the efficiency costs of selection, as emphasized by [Einav et al. \(2010\)](#). By abstracting from market power, we isolate the welfare consequences of information-driven selection from those arising from strategic pricing distortions. As a result, any welfare gains or losses can be attributed purely to selection effects and do not depend on the details of the pricing model.

In addition, we abstract from demand-side frictions and assume that consumers always choose the product with the lowest price. We also restrict our analysis to symmetric equilibria, in which Firms 1 and 2 offer identical prices prior to any merger. Firm 3 may behave differently, as it does not possess information about consumers and operates under a cost structure that differs from that of the other two firms. Finally,

we want to highlight that even with these assumptions, multiple equilibria may arise. In Appendix A, we provide a simple example illustrating how multiplicity can occur in this setting. For our analysis, we focus on the most efficient equilibrium—i.e., the equilibrium that minimizes the total cost of serving the market—as a benchmark for comparing market efficiency before and after the merger. Throughout this section, we focus on the merger between Firms 1 and 2 to illustrate the main mechanisms and to streamline the exposition.

2.1 Prices

We first analyze firms' pricing strategies before and after a merger between Firms 1 and 2. Before the merger, Firm $j \in \{1, 2\}$ offers two prices, p_l and p_h , corresponding to signals $s_j = l$ and $s_j = h$, respectively. Taking consumers' choices into account, Firm j recognizes that if a consumer accepts the high price p_h , it must be that both firms received signal h ; otherwise, the consumer would choose the firm offering the lower price p_l . As a result,

$$p_h = \mathbb{E}(c_1 | s_1 = h, D = 1) = \mathbb{E}(c_1 | s_1 = h, s_2 = h).$$

If a consumer accepts the low price, it implies that at least one of the signals from Firms 1 and 2 is l . Therefore,

$$p_l = \mathbb{E}(c_1 | s_1 = l, D = 1) = \mathbb{E}(c_1 | l \in \{s_1, s_2\}).$$

Firm 3, on the other hand, is fully uninformative and therefore offers a single uniform price, p_3 . We denote Firm 3's average cost of serving the entire market by $\mathbb{E}(c_3) = \mu_0 c_3^h + (1 - \mu_0) c_3^l$.

After Firms 1 and 2 merge, they observe both signals and offer prices based on three signal profiles, $\{l, l\}$, $\{h, l\}$, and $\{h, h\}$, separately.¹⁰ This essentially captures

¹⁰We rule out equilibria in which the merged firm charges the same price for signal profiles $\{l, l\}$ and $\{h, l\}$, as in the pre-merger case. Such equilibria are trivial and imply no change in market outcomes following the merger. Throughout, we assume that firms fully exploit the information available to them and set prices based on perceived risk.

the information synergies arising from the merger. Specifically,

$$\begin{aligned} p_{ll} &= \mathbb{E}(c_1 | s_1 = l, s_2 = l), \\ p_{hl} &= \mathbb{E}(c_1 | \{s_1, s_2\} = \{h, l\}), \\ p_{hh} &= \mathbb{E}(c_1 | s_1 = h, s_2 = h). \end{aligned}$$

If both signals are h , the probability that a consumer is a high-risk type is updated according to Bayes' rule as follows:

$$\mu_{hh} = \frac{\mu_0 * q^2}{\mu_0 * q^2 + (1 - \mu_0) * (1 - q)^2}.$$

The posterior beliefs corresponding to the other signal profiles can be derived analogously. Overall, a greater number of h signals implies a higher probability that the consumer is high-risk, which in turn leads to a higher expected cost for insurers.

Market equilibrium, both before and after the merger, depends on Firm 3's costs relative to those of the other two firms. For example, if Firm 3 has an extremely high cost, in the sense that its average cost exceeds Firms 1 and 2's expected cost conditional on two h signals, that is, $\mathbb{E}(c_3) > \mathbb{E}(c_1 | s_1 = h, s_2 = h)$, then Firms 1 and 2 will always dominate Firm 3, both before and after the merger. In this case, consumers do not switch across firms when the merger occurs, and there are no changes in market efficiency. At the other extreme, if Firm 3 has an extremely low cost, in the sense that its average cost is below Firms 1 and 2's expected cost conditional on two l signals, that is, $\mathbb{E}(c_3) < \mathbb{E}(c_1 | s_1 = l, s_2 = l)$, then Firm 3 will always dominate, and again there are no changes in efficiency following the merger. For the remainder of the analysis, we focus on the nontrivial case in which Firm 3's cost is comparable to that of Firms 1 and 2 and impose the following assumption on the firms' cost structures.

Assumption 1 (Nontrivial cost structure).

$$\mathbb{E}(c_1 | s_1 = l, s_2 = l) < \mathbb{E}(c_3) < \mathbb{E}(c_1 | s_1 = h, s_2 = h).$$

2.2 Efficiency

Next, we analyze how a merger between Firms 1 and 2 affects efficiency in this market. In our setup, insurance is mandatory, so we do not consider efficiency gains or losses arising from consumers being selected into or out of the market. Instead, we focus on *allocative efficiency*—that is, how consumers are sorted across firms so as to minimize the cost of serving them. We provide illustrative examples showing that information synergies arising from mergers can either decrease or increase market efficiency, then more generally characterize the conditions under which each scenario arises.

Efficiency Loss from Information Synergies In Figure 1, we present an example of a cost structure under which market efficiency decreases after Firms 1 and 2 merge. We first analyze the pre-merger scenario shown in Figure 1a. The vertical axis represents firms’ expected costs given their posterior belief μ that a consumer is high-risk, which is shown on the horizontal axis. The flatter line shows the expected cost function for Firms 1 and 2, while the steeper line shows that of Firm 3. The cost structure illustrates that Firm 3 has a much lower cost of serving low-risk consumers; however, its cost of serving high-risk consumers is similar to that of the other two firms. This difference implies that Firm 3 has a comparative advantage in serving low-risk consumers, while Firms 1 and 2 are comparatively better suited to serve high-risk consumers.

Given our assumption that firms charge prices equal to expected costs, Firm $j \in \{1, 2\}$ charges, prior to the merger, a high price p_h corresponding to the posterior belief given two h signals (i.e., μ_{hh}), and a low price p_l corresponding to the posterior belief given at least one l signal (i.e., $\mu_{l \in \{s_1, s_2\}}$). In contrast, Firm 3 is fully uninformative about a consumer’s risk and therefore sets its price equal to its expected cost conditional on being chosen. If Firm 3 serves the entire market, this conditional expectation equals the market average, so $p_3 = \mathbb{E}(c_3)$. When $\mathbb{E}(c_3) < p_l$ as shown in Figure 1a, Firm 3 undercuts Firms 1 and 2 and serves the entire market in equilibrium. This allocation is efficient, as Firm 3 has a lower cost of serving consumers.

We now analyze what happens after Firms 1 and 2 merge in Figure 1b. The key point is that the merged firm now receives two signals. As a result, the original posterior belief based on “at least one low signal” splits into two separate cases: μ_{ll}

and μ_{hl} , highlighted in red in Figure 1b. In our symmetric setup, note that when the signals are mixed—one low and one high—the posterior belief equals the prior μ_0 .

The merged firm can now identify a subset of particularly safe consumers, namely those with two low signals. Their posterior risk is even lower than before, which allows the merged firm to offer them a price p_{ll} below the market average price charged by Firm 3 prior to the merger (by Assumption 1). In this situation, the merged firm cream-skims the lowest-risk consumers. However, overall market efficiency declines, because Firm 3 has a lower cost of serving these consumers.

Once the merged firm starts serving the lowest-risk segment, Firm 3 must update its beliefs. It now recognizes that the remaining consumers are those with at least one high signal. Firm 3 therefore recalculates its expected cost conditional on at least one high signal, $\mathbb{E}(c_3|h \in \{s_1, s_2\})$. If this expected cost exceeds the price charged by the merged firm to consumers with mixed signals (p_{hl} , as highlighted in red in Figure 1b), Firm 3 also loses this group of consumers. As a result, market efficiency declines even further. When this occurs, Firm 3 updates its beliefs again, realizing that it is left only with consumers who have two h signals. If Firm 3's expected cost of serving this highest-risk segment is lower than the price charged by the merged firm given two h signals, Firm 3 ultimately serves the hh segment in equilibrium, while the merged firm serves the remaining lower-risk segments.

Figure 1 presents a scenario in which a merger between Firms 1 and 2 leads to a decrease in market efficiency. Post-merger, there is a clear misallocation of consumers across firms. Firm 3, which has a much lower cost of serving low-risk consumers, ends up serving the high-risk segment of the market after the merger. This outcome arises because the information synergies generated by the merger allow Firms 1 and 2 to cream-skim low-risk consumers, even though Firm 3 is the more efficient provider for that segment. The resulting misallocation increases the total cost of serving the market from a social planner's perspective.

Efficiency Gain from Information Synergies We consider an alternative cost structure in which a merger between Firms 1 and 2 increases market efficiency. In Figure 2, we plot the expected cost functions for all firms. The key difference relative to Figure 1 is that Firms 1 and 2 now have a lower cost of serving low-risk consumers, while Firm 3 has a lower cost of serving high-risk consumers. Equivalently, Firm 3's expected cost function is flatter than those of the other two firms. Under this

cost structure, comparative advantage is reversed: Firms 1 and 2 have a comparative advantage in serving low-risk consumers, whereas Firm 3 has a comparative advantage in serving high-risk consumers.


We again assume that Firm 3's market average price p_3 is lower than the low price p_l charged by Firms 1 and 2, so Firm 3 serves the entire market prior to the merger. This outcome, however, is not the most efficient allocation given the cost structure, because Firms 1 and 2 have a lower cost of serving low-risk consumers.

The merger between Firms 1 and 2 allows them to observe two signals, so the original posterior belief based on "at least one low signal" splits into two separate cases, μ_{ll} and μ_{hl} . The merged firm now offers the safest group of consumers a much lower price, p_{ll} , which is below Firm 3's market average price prior to the merger (by Assumption 1). When this occurs, the safest consumers switch to the merged firm after the merger. Given the cost structure shown in Figure 2, this reallocation of consumers to the merged firm increases market efficiency, since the merged firm has a lower expected cost of serving this group of consumers, that is, $\mathbb{E}(c_1|ll) < \mathbb{E}(c_3|ll)$.


Once the merged firm starts serving the lowest-risk segment, Firm 3 must update its beliefs. It now recognizes that the remaining consumers are those with at least one high signal. In Lemma 1 in Appendix A, we show that under the cost structure depicted in Figure 2, Firm 3's expected cost conditional on observing at least one high signal, $\mathbb{E}(c_3|h \in \{s_1, s_2\})$, is lower than the price charged by the merged firm to consumers with mixed signals, p_{hl} . As a result, Firm 3 sets its price at $p_3 = \mathbb{E}(c_3|h \in \{s_1, s_2\})$ and retains all consumers with at least one high signal, leaving efficiency for this group unchanged.

Comparing the two illustrative examples in Figures 1 and 2, we see that whether information synergies arising from mergers increase or decrease market efficiency depends critically on firms' cost structures. In Figure 1, the merged firm has a cost advantage in serving high-risk consumers, but the information advantage generated by the merger leads it to cherry-pick low-risk consumers, resulting in a reduction in efficiency. By contrast, in Figure 2, the merged firm has a cost advantage in serving low-risk consumers, and the information synergies generated by the merger allow it to select precisely these consumers. As a result, the merger is efficiency-improving in this case.

Importantly, the efficiency effects also depend on the magnitude of the information synergies. In our setting, this corresponds to the extent to which observing two signals



(a) Pre-merger



(b) Post-merger


Figure 2: Information synergies arising from the merger of Firms 1 and 2 lead to an increase in allocative efficiency. The horizontal axis shows posterior beliefs that a consumer is high-risk, and the vertical axis shows firms' expected costs; the steeper line corresponds to Firms 1 and 2 and the flatter line to Firm 3. Firm 3 has a relatively low cost compared to Firms 1 and 2. It serves the entire market prior to the merger, and consumers with at least one high signal after the merger.

shifts posterior beliefs relative to the belief based on observing at least one low signal. For example, in Figure 2b, if information synergies are sufficiently weak—so that the posterior belief μ_{ll} remains close to $\mu_{l \in \{s_1, s_2\}}$ and lies in a region where Firm 3’s expected cost of serving consumers with two low signals is still lower than that of the merged firm—then reallocating these consumers to the merged firm after the merger would instead reduce overall market efficiency.


The examples shown in Figures 1 and 2 both impose the assumption that Firm 3 has a relatively low cost compared to Firms 1 and 2, that is, $\mathbb{E}(c_3) < \mathbb{E}(c_1 | l \in \{s_1, s_2\})$. In this case, Firm 3 serves the entire market prior to the merger. We now consider an alternative scenario in which Firm 3 has a relatively high cost compared to Firms 1 and 2, that is, $\mathbb{E}(c_3) > \mathbb{E}(c_1 | l \in \{s_1, s_2\})$. In the pre-merger case shown in Figure 3a, suppose that Firm 3’s market average price, $\mathbb{E}(c_3)$, exceeds the low price p_l charged by Firms 1 and 2. In this case, Firms 1 and 2 attract consumers with at least one l signal, while Firm 3 serves consumers with two h signals. The price charged by Firm 3 is therefore $p_3 = \mathbb{E}(c_3 | hh)$.

After Firms 1 and 2 merge, as shown in Figure 3b, we again observe that the original posterior belief based on “at least one low signal” splits into two separate cases, μ_{ll} and μ_{hl} , because the merged firm can now observe two signals. In this scenario, the safest group of consumers are separated from those with mixed signals, which in turn creates incentives for Firm 3 to compete for the consumers with mixed signals. If Firm 3’s cost of serving the intermediate- and high-risk consumer groups is sufficiently low, it can set its price p_3 equal to its expected cost conditional on observing at least one high signal, $\mathbb{E}(c_3 | h \in \{s_1, s_2\})$. If this price is lower than the price charged by the merged firm for consumers with mixed signals, p_{hl} , as illustrated in Figure 3b, the merged firm loses consumers with mixed signals, and Firm 3 serves consumers with at least one h signal. As Firm 3, which can more efficiently serve low-risk consumers, gains a larger market share among lower-risk consumers, overall market efficiency increases.

Summary We formalize the effect of a merger between Firms 1 and 2 on market efficiency in the following proposition. We say that Firm 3 has a *relatively low cost* if $\mathbb{E}(c_3) < \mathbb{E}(c_1 | l \in \{s_1, s_2\})$, and a *relatively high cost* if $\mathbb{E}(c_3) > \mathbb{E}(c_1 | l \in \{s_1, s_2\})$. The cost structures in Figures 1 and 2 correspond to the case in which Firm 3 has a relatively low cost, while the cost structure in Figure 3 corresponds to the case in



(a) Pre-merger



(b) Post-merger

Figure 3: Information synergies arising from the merger of Firms 1 and 2 lead to an increase in allocative efficiency. The horizontal axis shows posterior beliefs that a consumer is high-risk, and the vertical axis shows firms' expected costs; the flatter line corresponds to Firms 1 and 2 and the steeper line to Firm 3. Firm 3 has a relatively high cost compared to Firms 1 and 2. It serves the hh segment prior to the merger, and consumers with at least one high signal after the merger.

which Firm 3 has a relatively high cost.

Proposition 1 (Efficiency).

1. *Suppose Firm 3 has a relatively low cost. A merger between Firms 1 and 2 strictly decreases market efficiency if $\mathbb{E}(c_1|s_1 = l, s_2 = l) > \mathbb{E}(c_3|s_1 = l, s_2 = l)$, and strictly increases market efficiency if $\mathbb{E}(c_1|s_1 = l, s_2 = l) < \mathbb{E}(c_3|s_1 = l, s_2 = l)$.*
2. *Suppose Firm 3 has a relatively high cost. A merger between Firms 1 and 2 weakly increases market efficiency. The increase is strict if and only if $\mathbb{E}(c_3|h \in \{s_1, s_2\}) \leq \mathbb{E}(c_1|\{s_1, s_2\} = \{l, h\})$.*

This proposition summarizes how a merger affects market efficiency under different cost structures. When Firm 3 has a relatively low cost, it serves the entire market prior to the merger. The information synergies arising from the merger allow the merged firm to cream-skin the low-risk segment of the market. Whether this reallocation increases or decreases market efficiency depends on whether shifting low-risk consumers from Firm 3 to the merged firm raises or lowers the cost of serving them. If the merged firm has a higher cost of serving these consumers, the reallocation reduces efficiency; otherwise, it increases efficiency. When Firm 3 has a relatively high cost, the information synergies generated by the merger between Firms 1 and 2 separate the lowest-risk consumers from the rest of the market, which creates an opportunity for Firm 3 to compete for the intermediate-risk segment. If Firm 3's cost of pooling these consumers with the highest-risk segment is sufficiently low, Firm 3 attracts these consumers, and the resulting reallocation is efficiency improving. We provide the formal proof of Proposition 1 in Appendix A.

2.3 Consumer Welfare

The merger between Firms 1 and 2 has heterogeneous effects across consumers. This section evaluates its impact on consumer welfare conditional on signal realizations rather than true risk types. Consumers are classified into three groups based on their signal profiles— ll , lh , and hh —which correspond to low-, intermediate-, and high-risk signal groups, respectively. It is important to note that each signal group contains both high- and low-risk types. However, the posterior probability that a consumer is

high-risk increases with the number of h signals received. For example, the posterior probability of being high-risk in the ll group is substantially lower than in the other two groups.

First, we observe that consumers in the low-risk signal group—those for whom two low signals are observed—unambiguously benefit from the increase in information precision of the merged firm, as they can now be more accurately identified and priced according to their true expected risk.

In contrast, consumers in intermediate-risk signal group—those for whom mixed signals are observed—are worse off. Prior to the merger, these consumers are pooled with the lowest-risk segment. After the merger, however, the increase in information precision allows the merged firm to separate them from the low-risk signal group. For example, in Figure 1b, consumers with mixed signals are served by the merged firm after the merger and pay a price p_{hl} , which is higher than the market-average price prior to the merger. Moreover, consumers in intermediate-risk signal group may instead be pooled with the high-risk segment after the merger. In Figures 2b and 3b, consumers with mixed signals are served by Firm 3 after the merger and are pooled with high-risk consumers, leading them to pay a higher price than in the pre-merger scenario.

The impact of information synergies on the welfare of the consumers in high-risk signal group is more subtle. In Figures 1a and 2a, consumers with two h signals are pooled with all other consumers prior to the merger, which allows them to pay a price well below the expected cost they impose on insurers. After the merger, however, the increase in the merged firm’s information precision leads to these high-risk consumers being separated out, and they face a substantial price increase, as illustrated by the green dashed line in Figures 1b and 2b.

Somewhat surprisingly, consumers in the high-risk signal group may nevertheless benefit from information synergies. When Firm 3 has a relatively high cost, as illustrated in Figure 3a, consumers with two h signals pay a substantially higher price than the rest of the market. After the merger, once the lowest-risk segment is separated out, Firm 3 may find it profitable to compete for consumers in the intermediate-risk segment. If Firm 3’s cost of pooling consumers with mixed signals or two h signals is sufficiently low to undercut the merged firm, then consumers in the high-risk signal group benefit from the resulting price reduction.

Taken together, these theoretical results highlight that information synergies from

mergers have important implications for market efficiency and can generate heterogeneous welfare effects across consumer risk types. The magnitude, and even the direction of these effects depends critically on the underlying cost structure and the strength of information synergies in a given market. In practice, these objects are inherently empirical and may vary substantially across settings. In the remainder of the paper, we quantify these effects in the Italian auto insurance market using merger simulations based on a flexible empirical framework of imperfect competition in selection markets described in Section 3.

3 Empirical Framework

To assess the quantitative importance of merger-driven information synergies, we study the Italian auto insurance market. Our analysis uses the data and empirical framework in [Cosconati et al. \(2025\)](#), which models competition among firms with heterogeneous information in this market. We first describe the key institutional features of the market and then briefly summarize the empirical model, focusing on how it extends the stylized framework in Section 2.

Italian Auto Insurance Market Our empirical setting is the motor vehicle liability insurance market in Italy. This mandatory insurance covers damages to third parties’ health and property in at-fault accidents. Policies are renewed annually and are exclusive, meaning that consumers may purchase coverage from only one insurer at a time. Insurers are also legally prohibited from rejecting applicants. This market is sizable: in 2018, approximately 31 million contracts were underwritten and more than 4 million claims were filed. Nationwide, around 50 firms compete in the provision of this coverage. In the empirical analysis, we focus on the ten largest firms and aggregate the remaining firms into a single “Firm 11” category representing the competitive fringe, for computational tractability.

The baseline contract in this market is relatively simple: only a negligible fraction of policies include a deductible, and all contracts are subject to mandatory minimum liability coverage of 1 million euros for property damage and 6 million euros for bodily injury. Consumers may also select optional contractual clauses, which can either increase or decrease the premium depending on whether they raise or reduce

the insurer’s risk exposure.¹¹ Premiums in this market are determined by insurers’ actuarial pricing algorithms, which incorporate a broad set of observable risk factors, including driver age, vehicle characteristics, place of residence, driving history, and optional contractual features. In addition, sales agents may grant discretionary discounts to individual consumers.

The data consist of a sample of liability insurance contracts sold between 2013–2021 in the province of Rome. For each consumer, we observe a rich set of risk characteristics, the insurer from which coverage is purchased, the premium paid, and the specific coverage and contractual clauses selected. The data are longitudinal and are matched to detailed information on each sampled consumer’s claim frequency and claim severity in every contract year.

Empirical Model The stylized model in Section 2 isolates the key mechanisms through which merger-driven information synergies affect market outcomes, but abstracts from several features of the empirical environment. The empirical framework in Cosconati et al. (2025) accommodates these features by allowing for imperfect competition, rich heterogeneity in firms’ costs and information precision, a continuous distribution of consumer risk, and flexible consumer demand.

First, the empirical framework allows for a substantially more flexible information structure than the stylized model. Consumers are indexed by a continuous risk type θ , which represents their expected cost to insurers, expressed in monetary terms. Although the realized risk type is not observed ex ante by insurers, the population distribution of θ is common knowledge. To better capture real-world risk-rating practices, in which insurers typically assign continuous risk scores to consumers, firm j is assumed to observe a continuous signal $\hat{\theta}_j \sim N(\theta, \sigma_j^2)$, rather than the binary signal in the stylized model. The parameter $1/\sigma_j$ captures the firm’s information precision. As in the stylized model, signals are private and independent across firms conditional on the true risk type θ . However, the signal distributions, and thus information precision, can vary flexibly across competing firms.

Second, the empirical model allows firms to set premiums strategically and compete in prices. Specifically, firm j sets the premium for a given consumer as a linear function of the expected cost, conditional on its signal realization and on being chosen

¹¹For example, the “Exclusive Driving” clause provides coverage only when the vehicle is driven by the individual named in the contract, thereby reducing the insurer’s risk exposure.

by the consumer:

$$p_j(\hat{\theta}_j) = \alpha_j + \beta_j \mathbb{E}(\theta | \hat{\theta}_j, D = j).$$

Here, α_j and β_j are pricing coefficients that firm j chooses optimally in equilibrium to maximize profits. The parameter α_j captures the baseline markup, while β_j governs the sensitivity of the premium to expected risk. In the stylized model, we impose $\alpha_j = 0$, and $\beta_j = 1$, which forces prices to equal expected costs and effectively rules out market power. In contrast, the empirical setting features imperfect competition, allowing firms to charge premiums above expected cost.

As in the stylized model, the expected cost conditional on a firm's signal and on the consumer selecting its contract, $\mathbb{E}(\theta | \hat{\theta}_j, D = j)$, is an equilibrium object, since it depends on all firms' pricing strategies and on consumer demand responses. Conditioning on the contract being chosen incorporates a winner's curse effect: firms recognize that winning the consumer may reveal adverse information about the underlying risk type.

Third, in addition to heterogeneity in firms' information structures, the empirical framework allows for heterogeneity in their cost structures. Specifically, the net profit that firm j derives from a consumer of type θ when the consumer purchases a contract from firm j is modeled as:

$$\text{Net Profit} = p_j(\hat{\theta}_j) - mc_j - k_j \theta,$$

where $p_j(\hat{\theta}_j)$ denotes the premium charged based on the firm's signal. The parameter mc_j captures firm-specific marginal servicing costs that are independent of risk, while k_j reflects the firm's efficiency in processing claims, with a lower k_j indicating higher cost efficiency.¹² The cost coefficients (mc_j, k_j) correspond to the intercept and slope, respectively, of the expected cost curve shown in Figures 1–3. A firm is more efficient at servicing high-risk consumers when its cost curve is flatter, that is, when it has a lower value of k_j .

Expected profits are obtained by integrating the firm's net profit over both the

¹²Cosconati et al. (2025) does not directly estimate mc_j . Instead, they estimate a firm-specific “net benefit” parameter c_j , which captures the expected future profits from customer retention or cross-selling opportunities, net of marginal servicing costs. The authors then use a back-of-the-envelope calculation, based directly on the data, to approximate these future benefits and thereby recover each firm's marginal cost mc_j . This approach preserves tractability by avoiding the need to explicitly model the dynamic trade-offs firms face when setting premiums.

distribution of signals and the population distribution of risk types, accounting for the probability that a consumer selects each firm. Firms then simultaneously choose their pricing coefficients (α_j, β_j) to maximize expected profits in a static game.

Finally, the empirical model allows for a more flexible demand system in which consumers respond to prices and have preferences over unobserved product attributes. Rather than assuming that consumers always choose the lowest-priced contract, consumer i 's indirect utility from purchasing a contract from firm j is given by

$$u_{ij} = -\gamma p_j(\hat{\theta}_j) + \xi_j + \varepsilon_{ij},$$

where γ denotes the price sensitivity parameter and ξ_j captures product-specific unobserved attributes, such as product quality or brand loyalty. These consumer preference parameters can vary with observable individual characteristics, such as age and geographic location. The term ε_{ij} represents idiosyncratic preference shocks, which is assumed to follow a Type I extreme value distribution.

Summary of Structural Estimates We now briefly summarize the main empirical patterns from the structural estimates in [Cosconati et al. \(2025\)](#). The demand estimates indicate that senior drivers are generally less price sensitive, whereas drivers located in larger urban centers are more price sensitive. Estimated preferences for unobserved product attributes are broadly similar across low- and high-risk groups. On the supply side, there is substantial heterogeneity across firms in both information precision and cost structures. Firms with the least advanced risk-rating technologies exhibit signal standard deviations more than 50 percent higher than those of firms with the highest precision in risk assessment. Firms also differ markedly in their efficiency in processing claims and in their risk-irrelevant marginal costs of serving consumers. Interestingly, firms with less accurate risk-rating algorithms tend to have more efficient cost structures, leading to distinct types of comparative advantages.

4 Merger Simulations

For the merger simulations, we use the structural estimates reported in [Cosconati et al. \(2025\)](#) to quantify the effects of information synergies. In particular, we simulate all possible mergers between any two of the ten largest insurers, yielding 45 distinct

merger pairs. To capture the information synergies generated by a merger, we assume that the merged firm observes both signals that were previously private information of the merging parties, as in the stylized model. Under the assumption that signals are independent (conditional on the true risk type) and normally distributed, observing two signals is therefore equivalent to receiving a combined signal with higher precision. The variance of this combined signal for the merged firm is given by

$$\sigma_m^2 = \frac{\sigma_j^2 \sigma_{j'}^2}{\sigma_j^2 + \sigma_{j'}^2},$$

where σ_j^2 and $\sigma_{j'}^2$ denote the variances of the signal distributions for merging firms j and j' , respectively. Because firms differ in their information precision, the magnitude of information synergies varies across merger scenarios.

To conduct merger simulations, we must also specify how cost structures and product attributes evolve post-merger. Our baseline specification assumes that the merged firm consolidates the insurance products of the two firms, and that its cost parameters and unobserved product attribute are given by the weighted averages of the corresponding pre-merger values, with weights proportional to their market shares.¹³ As an alternative, we also implement simulations under a more favorable scenario in which the merged firm inherits the lowest marginal cost and the highest claim processing efficiency among the merging parties. The two approaches yield qualitatively similar results. We therefore focus on the baseline specification, as it provides a neutral benchmark and avoids imposing strong assumptions about post-merger cost synergies.¹⁴

Computing the Post-Merger Equilibrium Solving for the post-merger equilibrium is computationally nontrivial. A merger changes not only market structure but also the information available to the merged firm, thereby altering firms' incentives

¹³The insurance products we study in this paper are standardized liability insurance and are close substitutes. [Mazzeo et al. \(2018\)](#) show that when products are close substitutes, a merged firm can reduce costs by eliminating duplicate products with similar characteristics. Such consolidation also occurs in practice; for example, after the merger between American Airlines and US Airways in 2013, the merged airline eliminated overlapping flights and redundant connections on routes where direct flights already existed ([Bontemps et al., 2022](#)).

¹⁴The merger simulation literature typically assumes unchanged cost structures following a merger (see, e.g., [Nevo \(2000\)](#)), or instead asks what level of cost efficiencies would offset the merger's effects on equilibrium outcomes.

to price consumers of different risk types. All firms must re-optimize their pricing strategies. These pricing responses affect consumer choices, which in turn change the composition of risks served by each insurer. The resulting changes in risk composition then alter firms' expectations about the consumers they attract, feeding back into their optimal pricing decisions. The post-merger equilibrium must therefore jointly determine firms' strategic pricing responses, consumer choices, and the endogenous sorting of risks across insurers.

To account for this feedback between pricing and consumer sorting, we solve for the post-merger equilibrium using the following iterative algorithm.

- **Initialize pricing strategies.** Start from a candidate set of pricing coefficients for all firms. In the merger simulations, we initialize the algorithm at the pre-merger equilibrium; for the merged firm, the initial coefficients are the market-share-weighted averages of the two merging firms' pre-merger coefficients.
- **Solve for consumer sorting given prices.** Holding the pricing coefficients fixed, compute the premiums implied by firms' current expectations of consumer risk and use these premiums to obtain consumer choice probabilities.
- **Update firms' risk assessments.** Use the resulting choice probabilities to update each firm's expected consumer risk conditional on its signal and on being chosen. Recompute premiums and choice probabilities, and iterate until these conditional risk expectations—and hence the sorting of consumers across firms—converge.
- **Update firms' pricing strategies.** Given the converged sorting pattern, evaluate the profit derivatives with respect to each firm's pricing coefficients and adjust those coefficients in the direction that increases profit.
- **Repeat until pricing converges.** Alternate between solving for consumer sorting and updating pricing strategies. To facilitate convergence, the algorithm reduces the adjustment step whenever a first-order condition changes sign across successive iterations.

Applying this procedure to each of the 45 merger scenarios, we re-solve the entire market equilibrium under the merged firm's new information, cost, and product characteristics, while holding the primitives of the non-merging firms fixed. This




Figure 4: Cream-skimming after mergers. The vertical axis plots the percentage change in the merged firm’s average claim payouts, while the horizontal axis shows its post-merger information precision. Each point represents a simulated merger between two of the ten largest firms (45 total mergers). The fitted line illustrates the negative relationship between information precision and average consumer risk.

allows us to capture both the direct effects of the merger and the endogenous equilibrium responses in firms’ pricing, consumer choices, and risk sorting. The resulting post-merger equilibria form the basis for the results reported below.

Information Advantages from Mergers The stylized model suggests that information synergies generated by a merger enable the merged firm to cream-skim low-risk consumers. Figure 4 illustrates how the average risk of the merged firm’s customers changes following the merger. The vertical axis reports the percentage change in the merged firm’s average claim payouts, while the horizontal axis shows its post-merger information precision. Each point represents a distinct merger scenario, yielding a total of 45 observations in the figure.

Figure 4 provides strong evidence that mergers generate substantial information advantages, enabling merged firms to cream-skim low-risk consumers. Across the 45 simulated merger scenarios, the merged firm’s average expected claim payouts decline by 7.2 to 15.5 percent relative to the pre-merger benchmark. These magnitudes are

economically significant and imply a sizable reallocation of low-risk consumers toward the merged firm.

Because the merging firms differ in their pre-merger information precision, the strength of information synergies varies across merger pairs. As a result, post-merger information precision exhibits substantial dispersion across the 45 scenarios, with the highest level of precision approximately 48 percent greater than the lowest. The figure further reveals a clear negative relationship between post-merger information precision and the percentage change in average claim payouts, as indicated by the downward-sloping fitted line. Mergers that result in higher information precision allow firms to attract a significantly lower-risk customer base.

Efficiency Effects We next examine how information synergies from mergers affect allocative efficiency in the market. Although improved information enables the merged firm to better identify and attract lower-risk consumers, it does not necessarily translate into lower aggregate costs of serving consumers in the market. In fact, our simulations show that information synergies can generate economically significant efficiency losses.

Figure 5 plots the percentage change in the average cost of serving the market following each merger. The vertical axis reports this change, while the horizontal axis displays the merged firm’s claim-processing cost parameter, k . It is clear to see from this figure that several merger scenarios—highlighted in the yellow region—lead to an increase in the market’s average cost, reflecting a deterioration in market efficiency.

When do efficiency losses arise? Figure 5 reveals a strong negative relationship between the merged firm’s claim-processing cost and overall market cost. When the merged firm has a low value of k , indicating high efficiency in processing claims, average market costs tend to rise after the merger. The mechanism is straightforward and consistent with the intuition highlighted in Figure 1 of the stylized model. The merged firm leverages its informational advantage to attract low-risk consumers, even in cases where it is particularly efficient at serving high-risk ones. From a social planner’s perspective, firms with low claim-processing costs should disproportionately serve high-risk consumers, since they can do so at lower cost. However, improved information allows the merged firm to cream-skim low-risk consumers instead. As a result, high-risk consumers are reallocated toward less efficient firms, leading to a misallocation of consumers across insurers. This distortion raises the average cost of




Figure 5: Information synergies and allocative efficiency. The figure plots the percentage change in the average cost of serving the market following each merger (vertical axis) against the merged firm's claim processing cost parameter k (horizontal axis). Each point represents a simulated merger scenario. Mergers in the highlighted yellow region result in higher average costs, indicating efficiency losses.

serving the market by up to 1.4 percent, or approximately 12 euros per contract.

While some mergers generate efficiency losses, information synergies can also produce efficiency gains in other scenarios, consistent with the predictions of the stylized model (see Figures 2 and 3). As shown in Figure 5, mergers between two insurers can reduce the market’s average servicing cost by up to 2.4 percent. These gains are concentrated in cases where the merged firm has a relatively high value of k , meaning it is comparatively more efficient at serving low-risk consumers. In such cases, improved information allows the merged firm to better align its cost advantage with the appropriate segment of consumers, thereby enhancing allocative efficiency rather than distorting it.

Competition authorities often rely on summary measures of market concentration, such as the Herfindahl–Hirschman Index (HHI), when evaluating mergers. To connect our analysis to this standard regulatory approach, Figure 6 plots the percentage change in the market’s average servicing cost (a measure of allocative efficiency) against the percentage change in HHI for each simulated merger scenario. The figure shows no systematic relationship between changes in market concentration and changes in market efficiency. Although the fitted line exhibits a slightly positive slope, the relationship is weak and statistically insignificant. In other words, increases in concentration do not reliably predict whether a merger improves or worsens allocative efficiency.

This finding suggests that summary concentration measures such as the HHI fail to capture important efficiency considerations in merger analysis. Even mergers that generate modest changes in HHI can produce meaningful efficiency losses, while others associated with larger concentration increases may improve efficiency. These results underscore the importance of incorporating firm-level information structures and cost heterogeneity into merger analysis, and more broadly, highlight the value of a comprehensive and systematic assessment framework.

Consumer Welfare Finally, we examine the welfare effects of mergers on consumers. Figures 7a and 7b report the percentage changes in consumer surplus for low-risk and high-risk consumers, respectively, where low-risk consumers are those below the median risk level and high-risk consumers are those above the median. In each figure, each block corresponds to a specific merger scenario—for example, the lower-left block represents the merger between Firms 1 and 2. Red shading indicates





Figure 6: HHI and allocative efficiency. The figure plots the percentage change in the average cost of serving the market following each merger (vertical axis) against the percentage changes in HHI (horizontal axis). Each point represents a simulated merger scenario. Mergers in the highlighted yellow region result in higher average costs, indicating efficiency losses.

an increase in consumer surplus, while blue shading indicates a decrease.


For low-risk consumers, we observe that welfare gains are generally modest and, in some merger scenarios, substantially negative. Across the 45 merger cases we consider, low-risk consumers experience an average decrease in consumer surplus of 0.16 percent. This pattern is consistent with the key mechanisms highlighted in our theoretical framework. As the merged firm improves its information precision, it can more accurately identify and price the safest group of consumers. However, consumers with intermediate risk—who were previously pooled with lower-risk types—are now separated out and face higher premiums. As a result, the aggregate welfare impact for below-median risk consumers is limited, and in some cases consumer surplus declines by more than 10 percent as shown in Figure 7a.

In contrast, the effects for high-risk consumers are markedly different. Figure 7b is dominated by red blocks, indicating that high-risk consumers tend to experience welfare gains following mergers. Across the 45 merger cases we consider, high-risk consumers experience an average increase in consumer surplus of 4.75 percent. Consistent with our theoretical predictions shown in Figure 3, improved information enables the merged firm to better separate risk types, which reshapes competitive pressures across segments. Firms with less precise information cannot profitably compete for the safest consumers, as doing so would expose them to adverse selection. They therefore redirect competition toward intermediate-risk consumers. In the post-merger equilibrium, high-risk consumers are pooled with these intermediate-risk consumers and face prices based on the average risk of the broader pool. This pooling generates cross-subsidization, lowering premiums for high-risk consumers and significantly increasing their welfare.

These findings highlight an important distributional dimension of information-driven mergers. Although average market outcomes may improve or deteriorate depending on the underlying cost structure, the welfare effects are clearly heterogeneous across consumer types. In particular, mergers can redistribute surplus across risk groups even when their aggregate impact appears modest. From a policy perspective, the desirability of such redistribution depends on the objectives of merger review. High-risk consumers may be disproportionately represented among economically vulnerable populations—for example, younger drivers, individuals in lower-income areas, or those with certain health conditions. In such cases, welfare gains for high-risk consumers may carry broader equity implications. Overall, our results suggest that



(a) Low-Risk Consumers



(b) High-Risk Consumers

Figure 7: Distributional effects of mergers on consumer surplus. Panel (a) plots the percentage change in consumer surplus for below-median (low-risk) consumers across simulated merger scenarios, while Panel (b) reports the corresponding changes for above-median (high-risk) consumers. Each block represents a merger between two of the ten largest firms. Red shading indicates an increase in consumer surplus, and blue shading indicates a decrease.

merger evaluations should consider not only aggregate efficiency effects, but also how information synergies reshape surplus across different groups of consumers.

5 Conclusion

Our paper develops a framework to analyze information synergies as a distinct channel through which mergers affect financial markets such as insurance and credit. We build a stylized model of insurer competition in which firms receive imperfect signals about consumers' risk types, and a merger allows them to integrate the signals held by both parties. Our theoretical results show that information synergies from mergers can increase the total cost of serving the market when the merged firm cream-skims low-risk consumers due to improved information precision, even though it is more efficient at serving high-risk consumers. We further show that information synergies can generate heterogeneous welfare effects across consumer risk types. Perhaps surprisingly, high-risk consumers may benefit in the post-merger equilibrium when they are cross-subsidized by intermediate-risk consumers.

We quantify the effects of information synergies from mergers on market efficiency and consumer welfare using merger simulations based on data from the Italian auto insurance market. We find that information synergies reduce the merged firm's average expected claim payouts by 7.2 to 15.5 percent. In a significant number of the merger scenarios we consider, however, the total cost of serving the market increases, by up to 1.4 percent, or approximately 12 euros per contract. Finally, consistent with our theoretical predictions, we find that high-risk consumers experience a significant increase in surplus, whereas the welfare gains for low-risk consumers are generally modest.

Our results have important implications for evaluating mergers among financial institutions, particularly in light of the growing importance of big data and AI-driven analytics. We highlight an information channel through which mergers can affect pricing, competition, and the allocation of surplus across consumers in financial markets. This channel is distinct from the conventional forces emphasized in merger analysis, such as changes in market concentration and cost efficiencies. Our findings suggest that merger assessments in financial markets should also account for how the consolidation of data and analytical capabilities changes firms' relative informational advantages. More broadly, when competitive concerns arise from the concentration of

information rather than traditional market power alone, our results point to a potential role for data-oriented remedies, such as data-access requirements, interoperability provisions, or restrictions on certain forms of data integration.

References

- ALLEN, J., R. CLARK, AND J.-F. HOUDE (2014): “The effect of mergers in search markets: Evidence from the Canadian mortgage industry,” *American Economic Review*, 104, 3365–3396.
- AN, Y. AND W. ZHAO (2019): “Dynamic efficiencies of the 1997 Boeing-McDonnell Douglas merger,” *The RAND Journal of Economics*, 50, 666–694.
- ARYAL, G., F. CILIBERTO, AND B. T. LEYDEN (2022): “Coordinated capacity reductions and public communication in the airline industry,” *The Review of Economic Studies*, 89, 3055–3084.
- ASHENFELTER, O. C., D. S. HOSKEN, AND M. C. WEINBERG (2015): “Efficiencies brewed: Pricing and consolidation in the US beer industry,” *The RAND Journal of Economics*, 46, 328–361.
- ASKER, J. AND V. NOCKE (2021): “Collusion, mergers, and related antitrust issues,” in *Handbook of Industrial Organization*, Elsevier, vol. 5, 177–279.
- ATER, I. AND O. RIGBI (2023): “Price transparency, media, and informative advertising,” *American Economic Journal: Microeconomics*, 15, 1–29.
- BANAL-ESTAÑOL, A. (2007): “Information-sharing implications of horizontal mergers,” *International Journal of Industrial Organization*, 25, 31–49.
- BERGER, A. N., R. S. DEMSETZ, AND P. E. STRAHAN (1999): “The consolidation of the financial services industry: Causes, consequences, and implications for the future,” *Journal of Banking & Finance*, 23, 135–194.
- BLICKLE, K., Z. HE, J. HUANG, AND C. PARLATORE (2025): “Information-based pricing in specialized lending,” *Journal of Financial Economics*, 172, 104135.
- BONTEMPS, C., K. REMMY, AND J. WEI (2022): “Ex-post evaluation of the American Airlines–US Airways merger: A structural approach,” *Journal of Transport Economics and Policy (JTEP)*, 56, 129–155.
- BORENSTEIN, S., J. KWOKA, AND L. WHITE (2004): “Rapid price communication and coordination: The airline tariff publishing case (1994),” *The Antitrust Revolution: Economics, Competition, and Policy*, 4.

- BRAGUINSKY, S., A. OHYAMA, T. OKAZAKI, AND C. SYVERSON (2015): “Acquisitions, productivity, and profitability: Evidence from the Japanese cotton spinning industry,” *American Economic Review*, 105, 2086–2119.
- BYRNE, D. P., N. DE ROOS, M. S. LEWIS, L. M. MARX, AND X. WU (2025): “Asymmetric information sharing in oligopoly: A natural experiment in retail gasoline,” *Journal of Political Economy*, 133, 2031–2088.
- CHEN, Z., C. CHOE, J. CONG, AND N. MATSUSHIMA (2022): “Data-driven mergers and personalization,” *The RAND Journal of Economics*, 53, 3–31.
- CHRISTENSEN, L. R. AND R. E. CAVES (1997): “Cheap talk and investment rivalry in the pulp and paper industry,” *The Journal of Industrial Economics*, 45, 47–73.
- COSCONATI, M., Y. XIN, F. WU, AND Y. JIN (2025): “Competing under information heterogeneity: Evidence from auto insurance,” *The Review of Economic Studies*, forthcoming.
- CRAWFORD, G. S., N. PAVANINI, AND F. SCHIVARDI (2018): “Asymmetric information and imperfect competition in lending markets,” *American Economic Review*, 108, 1659–1701.
- DE CORNIÈRE, A. AND G. TAYLOR (2024): “Data-driven mergers,” *Management Science*, 70, 6473–6482.
- DEYOUNG, R., D. D. EVANOFF, AND P. MOLYNEUX (2009): “Mergers and acquisitions of financial institutions: A review of the post-2000 literature,” *Journal of Financial Services Research*, 36, 87–110.
- DOYLE, M. P. AND C. M. SNYDER (1999): “Information sharing and competition in the motor vehicle industry,” *Journal of Political Economy*, 107, 1326–1364.
- EINAV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): “Estimating welfare in insurance markets using variation in prices,” *The Quarterly Journal of Economics*, 125, 877–921.
- ELIASON, P. J., B. HEEBSH, R. C. MCDEVITT, AND J. W. ROBERTS (2020): “How acquisitions affect firm behavior and performance: Evidence from the dialysis industry,” *The Quarterly Journal of Economics*, 135, 221–267.

- ELLIOTT, J. T., G. V. HOUNGBONON, M. IVALDI, AND P. T. SCOTT (2025): “Market structure, investment, and technical efficiencies in mobile telecommunications,” *Journal of Political Economy*, 133, 1401–1459.
- FARRELL, J. AND C. SHAPIRO (1990): “Horizontal mergers: An equilibrium analysis,” *The American Economic Review*, 107–126.
- FOCARELLI, D. AND F. PANETTA (2003): “Are mergers beneficial to consumers? Evidence from the market for bank deposits,” *American Economic Review*, 93, 1152–1172.
- GAL-OR, E. (1985): “Information sharing in oligopoly,” *Econometrica: Journal of the Econometric Society*, 329–343.
- GRIECO, P., J. PINKSE, AND M. SLADE (2018): “Brewed in North America: Mergers, marginal costs, and efficiency,” *International Journal of Industrial Organization*, 59, 24–65.
- HAUSWALD, R. AND R. MARQUEZ (2006): “Competition and strategic information acquisition in credit markets,” *The Review of Financial Studies*, 19, 967–1000.
- HE, Z., J. HUANG, AND J. ZHOU (2023): “Open banking: Credit market competition when borrowers own the data,” *Journal of Financial Economics*, 147, 449–474.
- HERRESTHAL, C., T. MAYS KAYA, AND A. NIKANDROVA (2023): “Data linkage between markets: Does the emergence of an informed insurer cause consumer harm?” *Working Paper*.
- JEZIORSKI, P. (2014): “Effects of mergers in two-sided markets: The US radio industry,” *American Economic Journal: Microeconomics*, 6, 35–73.
- KÜHN, K.-U. AND X. VIVES (1995): *Information exchanges among firms and their impact on competition*, Office for official publications of the European communities.
- LUCO, F. (2019): “Who benefits from information disclosure? The case of retail gasoline,” *American Economic Journal: Microeconomics*, 11, 277–305.
- MARSHALL, R. C., L. M. MARX, AND M. E. RAIFF (2008): “Cartel price announcements: The vitamins industry,” *International Journal of Industrial Organization*, 26, 762–802.

- MAZZEO, M. J., K. SEIM, AND M. VARELA (2018): “The welfare consequences of mergers with endogenous product choice,” *The Journal of Industrial Economics*, 66, 980–1016.
- MILLER, A. R. (2010): “Did the airline tariff publishing case reduce collusion?” *The Journal of Law and Economics*, 53, 569–586.
- NEVO, A. (2000): “Mergers with differentiated products: The case of the ready-to-eat cereal industry,” *The Rand Journal of Economics*, 395–421.
- PAGANO, M. AND T. JAPPELLI (1993): “Information sharing in credit markets,” *The Journal of Finance*, 48, 1693–1718.
- PANETTA, F., F. SCHIVARDI, AND M. SHUM (2009): “Do mergers improve information? Evidence from the loan market,” *Journal of Money, Credit and Banking*, 41, 673–709.
- RÖLLER, L.-H., J. STENNEK, AND F. VERBOVEN (2006): “Efficiency gains from mergers,” *European merger control: do we need an efficiency defence*, 84–201.
- ROSSI, F. AND P. K. CHINTAGUNTA (2016): “Price transparency and retail prices: Evidence from fuel price signs in the Italian highway system,” *Journal of Marketing Research*, 53, 407–423.
- SHAPIRO, C. (1986): “Exchange of cost information in oligopoly,” *The Review of Economic Studies*, 53, 433–446.
- WILLIAMSON, O. E. (1968): “Economies as an antitrust defense: The welfare trade-offs,” *The American Economic Review*, 58, 18–36.

A Appendix

A.1 Illustration of Multiple Equilibria

Consider an example of a cost structure shown in Figure 8. The vertical axis represents firms' expected costs given their posterior belief μ that a consumer is high-risk, which is shown on the horizontal axis. The flatter line shows the expected cost function for Firms 1 and 2, while the steeper line shows that of Firm 3. Given our assumption that firms charge prices equal to expected costs, Firm $j \in \{1, 2\}$ charges, prior to the merger, a high price p_h corresponding to the posterior belief given two h signals (i.e., μ_{hh}), and a low price p_l corresponding to the posterior belief given at least one l signal (i.e., $\mu_{l \in \{s_1, s_2\}}$).


Figure 8a illustrates one equilibrium. In this case, Firm 3, which is fully uninformative about a consumer's risk, sets its price p_3 equal to the market average, $\mathbb{E}(c_3)$. Suppose this price is lower than the low price p_l charged by Firms 1 and 2. Firm 3 then serves the entire market in equilibrium. This allocation is efficient, as Firm 3 has a lower cost of serving consumers.

Figure 8b illustrates another equilibrium. In this case, Firms 1 and 2 serve consumers who receive at least one l signal at price p_l , while Firm 3 serves the high-risk segment at price $\mathbb{E}(c_3|hh)$. This allocation is inefficient because Firm 3, despite having comparative advantages of servicing low-risk consumers, ends up serving only the high-risk segment. In our analysis, we focus on the efficient equilibrium as a benchmark for evaluating market efficiency before and after the merger.


A.2 Proofs

Proof of Proposition 1. Suppose that Firm 3 has a relatively low cost. In this scenario, before the merger, Firm 3 serves the entire market. After the merger, the merged firm can attract consumers with two l signals by the nontrivial cost assumption 1. We need to consider two cases.

Case 1. $\mathbb{E}(c_1|s_1 = l, s_2 = l) > \mathbb{E}(c_3|s_1 = l, s_2 = l)$. Then the shift of $\{s_1, s_2\} = \{l, l\}$ consumers strictly decreases market efficiency. Next, Firm 3 considers serving



(a) Equilibrium 1



(b) Equilibrium 2

Figure 8: Illustration of multiple equilibria. The horizontal axis shows posterior beliefs that a consumer is high-risk, and the vertical axis shows firms' expected costs; the flatter line corresponds to Firms 1 and 2 and the steeper line to Firm 3. Panel (a) illustrates an equilibrium in which Firm 3 charges $p_3 = \mathbb{E}(c_3)$ and serves the entire market. Panel (b) illustrates an equilibrium in which Firm 3 charges $p_3 = \mathbb{E}(c_3|hh)$ and serves only consumers with two h signals.

consumers who receive at least one h signal. If

$$\mathbb{E}(c_3|h \in \{s_1, s_2\}) < \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}),$$

then Firm 3 can retain all consumers who receive at least one h signal, and the result follows. Otherwise, consumers with $\{s_1, s_2\} = \{h, l\}$ also switch to the merged firm. Then all consumers who receive at least one l are served by the merged firm, which strictly decreases market efficiency due to

$$\mathbb{E}(c_3|l \in \{s_1, s_2\}) < \mathbb{E}(c_3) < \mathbb{E}(c_1|l \in \{s_1, s_2\}),$$

where the last inequality holds as Firm 3 has a relatively low cost. If

$$\mathbb{E}(c_3|\{s_1, s_2\} = \{h, h\}) \leq \mathbb{E}(c_1|\{s_1, s_2\} = \{h, h\}),$$

then the result follows immediately. Otherwise, consumers who receive two h signals also switch to the merged firm, implying that all consumers are served by the merged firm. This entire market movement strictly decreases efficiency since

$$\mathbb{E}(c_3) < \mathbb{E}(c_1|l \in \{s_1, s_2\}) < \mathbb{E}(c_1).$$

Case 2. $\mathbb{E}(c_1|s_1 = l, s_2 = l) < \mathbb{E}(c_3|s_1 = l, s_2 = l)$. Then the reallocation of $\{s_1, s_2\} = \{l, l\}$ consumers strictly increases market efficiency. Next, Firm 3 considers serving consumers who receive at least one h signal. Lemma 1 shows that Firm 3 can retain all such consumers, completing the argument.

Suppose that Firm 3 has a relatively high cost. Then before the merger, Firm 1 and 2 serves consumers with $l \in \{s_1, s_2\}$. Firm 3 serves consumers with $\{s_1, s_2\} = \{h, h\}$ if and only if

$$\mathbb{E}(c_3|\{s_1, s_2\} = \{h, h\}) < \mathbb{E}(c_1|\{s_1, s_2\} = \{h, h\}).$$

After the merger, the merged firm serves consumers with $\{s_1, s_2\} = \{l, l\}$. Firm 3 can attract all consumers who receive at least one h signal if and only if

$$\mathbb{E}(c_3|h \in \{s_1, s_2\}) \leq \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}).$$

If this condition fails, then consumers do not move and efficiency remain unchanged.

Now suppose this condition holds. If Firm 3 previously serves consumers with $\{s_1, s_2\} = \{h, h\}$, then only consumers with $\{s_1, s_2\} = \{h, l\}$ move to Firm 3, which strictly improves efficiency as

$$\mathbb{E}(c_3|\{s_1, s_2\} = \{h, l\}) < \mathbb{E}(c_3|h \in \{s_1, s_2\}) \leq \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}).$$

If Firm 3 previously does not serve consumers with $\{s_1, s_2\} = \{h, h\}$, then all consumers with $h \in \{s_1, s_2\}$ switch to Firm 3, which strictly improves efficiency since

$$\mathbb{E}(c_3|h \in \{s_1, s_2\}) \leq \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}) < \mathbb{E}(c_1|h \in \{s_1, s_2\}).$$

□

Lemma 1. *Suppose that Firm 3 has a relatively low cost. If*

$$\mathbb{E}(c_1|s_1 = l, s_2 = l) < \mathbb{E}(c_3|s_1 = l, s_2 = l),$$

then

$$\mathbb{E}(c_3|h \in \{s_1, s_2\}) < \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}). \quad (1)$$

That is, after the merger, Firm 3 can retain all consumers who receive at least one h signal.

Proof of Lemma 1. We first lay out the primitives of the model: signal precision q , population prior $\mu_h + \mu_l = 1$ (with $\mu_h = \mu_0$, and $\mu_l = 1 - \mu_0$), and cost parameters c_1^l , c_1^h , c_3^l , c_3^h . Without loss of generality, we can normalize $c_1^l = 0$ by adding a constant to all cost parameters. The proof shall proceed as follows: for each fixed primitive (q, μ_h, c_1^h) , we shall maximize the left-hand side of Inequality (1) and then show the maximum is below the right-hand side of Inequality (1). That is, for all q, μ_h, c_1^h ,

$$\max_{c_3^h, c_3^l} \mathbb{E}(c_3|h \in \{s_1, s_2\}) < \mathbb{E}(c_1|\{s_1, s_2\} = \{h, l\}).$$

By Bayes' rule, we can compute the posterior belief of being a high type given each signal profile.

$$\mu(\{s_1, s_2\} = \{l, l\}) = \frac{\mu_h(1-q)^2}{\mu_h(1-q)^2 + \mu_l q^2},$$

$$\begin{aligned}
\mu(l \in \{s_1, s_2\}) &= \frac{\mu_h(1 - q^2)}{\mu_h(1 - q^2) + \mu_l[1 - (1 - q)^2]}, \\
\mu(\{s_1, s_2\} = \{h, l\}) &= \mu_h, \\
\mu(\{s_1, s_2\} = \{h, h\}) &= \frac{\mu_h q^2}{\mu_h q^2 + \mu_l(1 - q)^2}, \\
\mu(h \in \{s_1, s_2\}) &= \frac{\mu_h[1 - (1 - q)^2]}{\mu_h[1 - (1 - q)^2] + \mu_l[1 - q^2]}.
\end{aligned}$$

We can then rewrite the left-hand side of (1) as follows:

$$\begin{aligned}
&\mathbb{E}(c_3 | h \in \{s_1, s_2\}) \\
&= \mathbb{E}(c_3) + [\mu(h \in \{s_1, s_2\}) - \mu_h](c_3^h - c_3^l) \\
&= \mathbb{E}(c_3) + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mathbb{E}(c_3) - \mathbb{E}(c_3 | \{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} \\
&< \mathbb{E}(c_1 | l \in \{s_1, s_2\}) + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mathbb{E}(c_1 | l \in \{s_1, s_2\}) - \mathbb{E}(c_3 | \{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} \\
&< \mathbb{E}(c_1 | l \in \{s_1, s_2\}) + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mathbb{E}(c_1 | l \in \{s_1, s_2\}) - \mathbb{E}(c_1 | \{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} \\
&= \mu(l \in \{s_1, s_2\})c_1^h + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mu(l \in \{s_1, s_2\}) - \mu(\{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} c_1^h
\end{aligned}$$

where the first inequality holds since Firm 3 has a relatively low cost; the second inequality holds since

$$\mathbb{E}(c_1 | s_1 = l, s_2 = l) < \mathbb{E}(c_3 | s_1 = l, s_2 = l).$$

Therefore, to prove Inequality (1), it suffices to show

$$\begin{aligned}
&\mu(l \in \{s_1, s_2\})c_1^h + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mu(l \in \{s_1, s_2\}) - \mu(\{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} c_1^h \\
&< \mu(\{s_1, s_2\} = \{h, l\})c_1^h,
\end{aligned}$$

which simplifies to

$$\mu(l \in \{s_1, s_2\}) + [\mu(h \in \{s_1, s_2\}) - \mu_h] \frac{\mu(l \in \{s_1, s_2\}) - \mu(\{s_1, s_2\} = \{l, l\})}{\mu_h - \mu(\{s_1, s_2\} = \{l, l\})} - \mu_h < 0.$$

Note that this inequality depends only on two parameters, the prior μ_h and the




Figure 9: The left-hand side of Inequality (1) as a function of μ_h and q .

precision q , as c_1^h cancels out. We plot the left-hand side of the inequality in Figure 9. As shown, the inequality holds across the entire parameter space $\mu_h \in (0, 1)$, and $q \in (0.5, 1)$.

□