University of Bergamo

Department of Management, Economics and Quantitative Methods

Date: September 30, 2019

I HEREBY RECOMMEND THAT THE THESIS PREPARED UNDER MY SUPERVISION BY

Michele Bisceglia

ENTITLED

Three Essays in Industrial Organization

BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN

ANALYTICS FOR ECONOMICS AND BUSINESS

Thesis Advisor
Prof. Sergio Ortobelli Lozza

Three Essays in Industrial Organization

By

Michele Bisceglia

Dissertation

Submitted in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy
in Analytics for Economics and Business
in the Department of Management, Economics and Quantitative Methods
at University of Bergamo, 2019

Bergamo, Italy

The sciences do not try to explain, they hardly even try to interpret, they mainly make models. By a model is meant a mathematical construct which, with the addition of certain verbal interpretations, describes observed phenomena. The justification of such a mathematical construct is solely and precisely that it is expected to work — that is, correctly to describe phenomena from a reasonably wide area.

John Von Neumann

Acknowledgements

It is incredibly difficult to express in a few lines my gratitude to all the wonderful people I met during the last three years.
First, I would like to thank my advisor, Sergio Ortobelli Lozza, because, whenever I needed help, I always could count on him. Next, I would like to thank the coauthors of the papers contained in this thesis: Salvatore Piccolo, who I can't thank enough for everything he did for me; Emanuele Tarantino, Jorge Padilla and Marco Pagnozzi, also for having been great hosts during my periods as a visiting PhD student at the University of Mannheim, at Compass Lexecon in Madrid and at the University of Naples Federico II, respectively. I am indebted to the coauthors of other ongoing research projects: Roberto Cellini, Luca Grilli, Heiko Karle, Francesco Menoncin, Luigi Siciliani and Odd Rune Straume, because from all of them I learned a lot. It is not possible to mention here all the professors I had the privilege to meet at the University of Bergamo and at various schools, seminars and conferences, who gave me incredibly valuable advice during these years. However, I must mention Lucia Maddalena, who first encouraged me to do research when I was a master's student at the University of Foggia.
For the great time I had during these years in Bergamo, I have to thank my colleagues and friends: Marco, Filippo and Lorenzo, and all the other office mates: Mehdi, Gabri, Matteo, Rossana, Iris, Laura, Giacomo, Michele and many others.
I would like to thank Ale, Yase, Yihan, Ele, Mattia, Federico, Cristina and the other friends who made me feel at home during my first experience abroad in Mannheim, and then Francesco, Niccolò, Leo, Raffaele, Vittorio and many others for the great time I had in Naples.
Last but not least, I am really grateful to my family, for always supporting me, especially when I felt guilty for being far from home in some hard times we had during these years, and my old friends (who are too many to mention them all), because, even tough the time we can spend together is much less than some years ago, they are like brothers to me.

Three Essays in Industrial Organization

Michele Bisceglia

Department of Management, Economics and Quantitative Methods
University of Bergamo
Bergamo, Italy
2019

ABSTRACT

This thesis contains three theoretical contributions in the field of industrial organization. The first chapter is concerned with the merger review process, and in particular with the role played by a professional advisor, who is hired by the merging parties to provide evidence about the *efficiencies* of the merger. It is shown that consumers are not necessarily better off when the advisor's contract is disclosed to the Antitrust Authority, due to a negative effect which hinges on a free-riding problem between advisor and authority in the information acquisition game, and is more relevant in highly competitive industries. The second chapter deals with the welfare effects of *platform parity agreements*, namely contractual provisions according to which a seller cannot charge different prices for the same product distributed through different platforms. The main result is that, differently from similar clauses examined in the available literature, in industries with a more complex vertical structure these provisions are likely to be pro-competitive even in the absence of efficiencies. Finally, a model of repeated adverse selection with spot contracts and persistent agents' type is considered in the third chapter, with the aim of investigating the impact of competition among managerial firms on the *ratchet effect*. Specifically, it is shown that, in a certain region of the parameters, a new kind of semi-separating equilibrium emerges, in which principals, as well as agents, randomize, the industry output is decreasing over time, and principals benefit from facing a more severe adverse selection problem.

Table of Contents

Introduction

Industrial organization (IO) can be broadly defined as the branch of microeconomics which studies the functioning of markets. Thus, IO analyzes the structure and behavior of firms (in terms of market strategy and internal organization) and draws normative implications (in terms of the appropriate government interventions aimed at improving social welfare) — see, e.g., Tirole (1988).

This thesis contains three essays which contribute to different strands of the theoretical IO literature. The first chapter, taken from a joint work with Prof. Salvatore Piccolo and Prof. Emanuele Tarantino, contributes to the literature on horizontal mergers. The basic trade off is as follows (e.g., Williamson, 1968): the merger increases the market power of the merged entity, but it can also bring efficiencies, in the form of a reduction of production costs. Specifically, our work relates to previous contributions analyzing the merger review process, which, in the available literature, is formalized as a game between the merging firms and the Antitrust Authority (henceforth, AA) which is in charge of the merger review process, under the assumptions that the proponent firms have an informational advantage *vis-à-vis* the AA as far as it concerns the efficiencies generated by the merger (e.g., Besanko and Spulber, 1993). Motivated by real world evidence, we examine the role played by a professional advisor, hired by the merging parties, in the merger review process. Specifically, by incurring an information acquisition cost, the advisor can gather information about the efficiencies and, if the information acquisition process is fruitful, he can certify these efficiencies *vis-à-vis* the AA, which, in turn, can conduct internal investigations to gather information about the efficiencies. Within this framework, we examine the effects on the accuracy of the merger review process (hence, on consumer surplus) of the disclosure of the advisor's contract, showing that the optimal *disclosure regime* depends on parameters related to the information acquisition process as well as to the market structure. In particular, it turns out that the disclosure of the advisor's contract harms consumers if and only if the cost of acquiring information is sufficiently low, and the region of parameters in which disclosure harms consumers expands as the market becomes less concentrated. Our results are shaped by two driving forces, namely a moral hazard problem in the contractual relation between the merging firms and the advisor and a free riding problem between the advisor and the AA in the information acquisition game.

The second chapter, taken from a joint work with Dr. Jorge Padilla and Prof. Salvatore Piccolo, covers another topic in the antitrust economics, by examining the pro- and the anti-competitive effects of a peculiar kind of Most Favoured Nation clause (henceforth, MFN) which is often observed in industries characterized by a rather complicated vertical structure. Specifically, we consider a monopolist selling its product through its direct sale channel and two competing platforms, which are in turn accessed by intermediaries competing for final consumers. In line with industry practices, we consider linear pricing and an agency business model, in which the monopolist sets the *access price* at which intermediaries relying on different platforms can buy the product, and platforms charge per-unit commissions to the monopolist. The MFN clause under investigation, known as *platform parity agreement*, requires the monopolist to post the same price on the two platforms. Thus, differently from the available models (e.g., Boik and Corts, 2016; Johansen and Vergé, 2016), we analyze a three-layer, rather than a simpler two-layer, industry, and a different kind of MFN, which does not impose any restriction to the price set on the direct channel. Within this setting, which is consistent with the airline ticket distribution industry, it is shown that the introduction of the parity provision, on the one hand, increases platforms' commissions (whereby leading the monopolist to lower its price on the relatively cheaper direct channel) but, on the other hand, is used by the monopolist as a commitment device to mitigate the multiple marginalization problem. As a result, in sharp contrast with the conclusions reached by the previous literature, we find that, even in the presence of an upstream monopolist and abstracting from efficiencies brought by the platforms, the considered parity agreement is pro-competitive, provided that the products sold by the monopolist through the different channels are perceived by consumers as rather homogeneous.

The third chapter, taken from a joint work with Prof. Marco Pagnozzi and Prof. Salvatore Piccolo, deals with competition among managerial firms in a dynamic contracting model. Specifically, we study a two-period economy populated by a continuum of perfectly competitive firms, each composed by a principal and an exclusive agent privately informed about his (persistent) production cost. Principals lack commitment power and can only use spot contracts. We show that the interplay between market competition and the *ratchet effect* brings out novel interesting equilibria of the game, yielding completely different implications in terms of industry dynamics as compared to the ones identified in the previous literature (e.g., Hart and Tirole, 1988). In particular, if players do not discount future profits, when the adverse selection problem is sufficiently severe, there exists an equilibrium in which both principals and agents randomize in the first period. In this equilibrium, principals' expected profit is an increasing function of the severity of the adverse selection problem, and the aggregate output is decreasing over time (i.e., we observe a *declining industry*) if the adverse selection problem is not too severe. If, on the contrary, players discount future profits, then, for intermediate values of the adverse selection problem, the game features a novel type of semi-separating equilibrium in which principals, rather than agents, randomize between the pooling and the separating contracts in both periods. Moreover, in this equilibrium, aggregate production is the same in both periods and principals' profits increase as the adverse selection problem worsens.

Chapter 1

M&A Advisory and the Merger Review Process

1.1 Introduction

M&A advisory fees from completed transactions totalled about 40 billion US dollars in 2017.[1] Goldman Sachs alone recorded total fees of about 3 billion US dollars from M&A deals. CEOs ask for advice to these firms, seeking for specialized help on decisions that are critical for the success of the transaction, like the identification of synergies and the design of the new company's structure after the deal is closed.[2] Advisors then support clients throughout the process: from the choice of the target to the execution of the due diligence, which may require reporting the transaction to the Antitrust Authority (AA), and disclosing to the Securities and Exchange Commission (SEC) the parties to be compensated and the relative arrangements.

Interestingly, recent legal developments highlight an increasing trend towards more extensive disclosure of advisory fees in M&A deals, including the circumstances in which these fees are payable.[3] Similarly, recent decisions of the Delaware Chancery Court have emphasized the need for greater transparency on advisors' potential conflicts of interest. The worry is that advisors may have perverse incentives to complete the transaction, possibly in opposition to the interests of shareholders.[4]

[1] Thomson Reuters and Freeman Consulting at www.thomsonreuters.co.jp/content/dam/openweb/documents/pdf/japan/market-review/2017/ma-4q-2017-e.pdf.

[2] The empirical literature is split on whether advisors provide valuable advice in M&A transactions. On the one hand, Bao and Edmans (2011) show that they have a positive impact on M&A returns, and Golubov et al. (2012) find evidence suggesting that top advisors are able to identify and structure mergers with higher synergies. On the other hand, Servaes and Zenner (1996) find no benefit of hiring an advisor, and the evidence in Rau (2000) is consistent with a negative impact of experts on post-merger returns.

[3] "Developments in Disclosure of Financial Advisor Fees in M&A Transactions", Detchert LLP, 2017; and "New Trends in M&A Transactions: to Disclose or Not to Disclose?", GS2 Law, 2011.

[4] Advisors need not be banks. There are many consulting companies that offer advice in mergers and acquisitions (e.g., McKinsey and International M&A Partnership, among many others). There are also examples in which advisors are industry 'veterans'. For example, Sumit Banerjee (a cement hotshot for over a decade) has become the reference for funds and strategic players whenever a new target comes into

Despite it is common for firms to use advisors in M&A transactions, and the greater scrutiny these advisors are subject to, it is unfortunate that little is known about how their presence impacts the success of M&A. What is the relationship between the accuracy of the merger evaluation process, the role of external advisors, and the rules governing the transparency of these deals? Do firms proposing a merger have an incentive to rely on external advice in order to support their proposal? Does the presence of experts alter the process of investigation on the welfare effects of the merger? Should the proponent firms be obliged to disclose advisory fees? This paper proposes a game-theoretic model to answer these questions.

We study a standard static framework à la Farrell and Shapiro (1990) in which two out of N symmetric firms, competing in a Cournot industry, propose a merger to an AA who adopts a consumer surplus standard, as is (roughly) the case in both the U.S. and EU antitrust law.[5] We assume that the basic trade-off is the standard one: the merger may increase market power but also create efficiencies (Williamson, 1968).[6] In the model, the proponent firms are uninformed about the (uncertain) efficiencies of the merger, but can hire an expert to gather information on their behalf. Firms cannot verify the information reported by the expert to the AA — i.e., they lack the technical skills necessary to assess whether such evidence is informative or not (the AA, by contrast, owns these skills). This asymmetry of information creates a moral hazard problem that can be solved by offering the expert an incentive compatible contract that pays a reward (known as *contingent fee*) only if the merger is approved. The AA is also uninformed about the efficiency improvement of the merger, but can run a costly internal investigation in order to learn its realized value. Hence, the model differs from Besanko and Spulber (1993) and Nocke and Whinston (2013) because firms are not privately informed about the merger characteristics *vis-à-vis* the AA, but both can acquire this information by engaging in a costly information acquisition activity. To make the problem interesting, we follow Besanko and Spulber (1993) in assuming that the merger harms consumer surplus in expected terms — i.e., behind the veil of ignorance the AA always rejects it.

the salemarket in India's heavily-crowded cement sector ('ACC confirms exit of Sumit Banerjee', The Economic Times, 2010).

[5]The U.S. Horizontal Merger Guidelines, issued in 2010 by the Department of Justice and the Federal Trade Commission, state that "the Agency will not challenge a merger if cognizable efficiencies ... likely would be sufficient to reverse the merger's potential harm to consumers in the relevant market." By the same token, the European Commission Guidelines on the Assessment of Horizontal Mergers under the Council Regulation on the Control of Concentrations between Undertakings, issued in 2004, state that "the relevant benchmark in assessing efficiency claims is that consumers will not be worse off as a result of the merger."

[6]Efficiency improvements can be measured not only in terms of production costs but also in terms of operational synergies (see, e.g., Gupta and Gerchak, 2002) or dynamic efficiencies and investments (Motta and Tarantino, 2018).

Within this framework, we analyze the welfare effects of the disclosure of the expert's contract as part of firms' due diligence. Our main findings follow. First, we show that hiring an advisor does not necessarily improve the chances that a merger is approved. Second, consumers are not necessarily better off when the advisor's contract is disclosed. These results hinge on the negative effect produced by a free-riding problem that arises between the expert and the AA in the information acquisition game. As we explain below, this problem is particularly relevant in highly competitive industries — i.e., in industries in which there are (relatively) more competitors.

Essentially, since gathering information about the merger's efficiency improvement is costly, both for the expert and the AA, when the cost of acquiring information is sufficiently large, both have an incentive to free ride on each other. On the one hand, the AA would like the expert to gather information so to base the decision regarding the merger only on the evidence that he collects, while saving on its own information acquisition cost. On the other hand, the expert would like the AA to collect information so to avoid paying the investigation cost, while being rewarded in case the merger is approved. In the region of parameters where this free-riding problem is relevant, disclosure plays a key role for consumer surplus.

Accordingly, the game unfolds differently depending on whether the AA is informed or not about the advisor's contract. Without disclosure the game is *de facto* simultaneous, and has the AA and the expert run into a coordination problem that may lead parties to play mixed strategies in equilibrium — i.e., the merging firms randomize between hiring and not hiring the expert (who, if hired, is induced to gather information), while the AA randomizes between collecting information and remaining uninformed. In this equilibrium both players collect information and, of course, the accuracy of the AA's decision increases if both choose to acquire information with a (relatively) high probability.

With disclosure, the structure of the game becomes sequential. However, for intermediate values of the contingent fee, the AA and the expert may play mixed strategies in the information acquisition game. If this is the case, then the firms can hire the expert optimally choosing a contingent fee in order to maximize their expected profit. It turns out that this constitutes the optimal choice of the firms when the cost of acquiring information is relatively low. On the contrary, for higher values of the information acquisition cost, the proponent firms exploit their first-mover advantage by committing not to hire an expert,[7] thus forcing the AA to acquire information (provided that the cost of doing so is not prohibitively high). Hence, in equilibrium, either the AA acquires information or the merger is rejected with certainty even though it may have been beneficial to consumers.

[7]Note that this is equivalent to submitting an uninformative report in case the AA asks for additional information during the merger review.

Avoiding this risk induces the AA to bear alone the cost of being informed.
Comparing the expected consumer surplus in these two regimes, we show that the disclosure of the expert's contract harms consumers if and only if the cost of acquiring information is sufficiently low. In the region of parameters in which mixed strategies are played in both disclosure regimes, the expert gathers information with the same probability regardless of his contract being disclosed, whereas the AA acquires information with higher probability when the expert's contract is not disclosed. Therefore, the disclosure of the expert's contract unambiguously harms consumers. Next consider the region of parameters in which with disclosure only the AA gathers information and becomes informed with some probability. In this case, when the cost of acquiring information is relatively low, without disclosure, the probability that the AA takes an informed decision is sufficiently large since both players acquire information with high probability, hence this regime is preferable from a consumer surplus viewpoint. By contrast, when the cost of acquiring information is sufficiently large, the probability that the AA takes an informed decision drops in the regime without disclosure and in that case enforcing transparency on the expert's contract benefits consumers.
Interestingly, the region of parameters in which disclosure harms consumers expands as the market becomes less concentrated. The reason is as follows. When the number of firms in the industry increases, the anticompetitive effect of the merger weakens. Hence, conditional on being informed, the AA rejects the merger with lower probability. Other things being equal, a lower probability of rejection tends to increase consumer surplus more without disclosure than with disclosure, because in the former regime both players acquire information (although they do so with probability less than one) while in the latter only the AA acquires information (with certainty).
Summing up, our model and results not only bring a novel theoretical twist, but they also provide useful insights to better understand how merger review protocols should take into account information that is not directly inherent to the structure of the industry (e.g., demand and market shares) and the particular deal under evaluation(e.g., costs, efficiency improvement, etc.).

The article proceeds as follows. After reviewing the related literature, we set up the model in Section 1.2 and first develop the analysis for the regime with disclosure (Section 1.3.2) and then the regime without disclosure (Section 1.3.3). In Section 1.3.4, we study the effect of disclosure on consumer surplus. Section 1.3.5 discusses some extensions and implications of the model. Section 1.4 concludes. Proofs are in the Appendix.

Related literature. Inspired by Williamson (1968), the modern approach to horizontal mergers has been developed by Farrell and Shapiro (1990), who formalize the basic welfare trade-off arising from horizontal mergers in a model with Cournot competition and increasing marginal costs (see, e.g., also McAfee and Williams, 1992).[8] These models assume complete information and a passive role for the AA.

Besanko and Spulber (1993) were the first to examine the role of the AA in a model with asymmetric information. They assume that merging firms are more informed than the AA about the merger's efficiency gains and show that, by committing to a conservative consumer surplus standard, the AA may increase aggregate surplus.[9] A different form of asymmetric information is considered in Armstrong and Vickers (2010), who determine the optimal permission set in a model in which the AA can verify the characteristics of the received merger proposal, but does not know which other merger projects were available to the firms (see, e.g., also Lyons, 2003). Building on the same idea, Nocke and Whinston (2013) analyze the optimal merger approval policy in a model with a single acquirer who can make a merger proposal to one of several heterogeneous firms. In particular, they assume that the AA observes the characteristics of the proposed merger only, which is the result of a bargaining process among the firms, and find that the AA optimally commits to a policy that imposes tougher standards on mergers involving firms with a larger pre-merger market share.

Finally, closest to us is Sørgard (2009), who studies a framework where the AA sets an activity level (how many mergers to investigate). He shows that an active merger control can lead to an adverse selection of proposed mergers, which in turn has implications for the possible decision errors made on investigated mergers. The optimal activity level trades off the possible mistakes from enforcement with the gains from deterring a potentially harmful merger.[10]

None of these papers studies the role of experts and explicitly models the information acquisition process.[11] By introducing both these ingredients, our model delivers policy implications on the relationship between the accuracy of the merger evaluation process, the role of external advisors and the rules governing how much transparent their deals should be.

[8]Relatedly, Deneckere and Davidson (1985) study mergers in a model with price competition and constant marginal costs, so when insiders raise prices, outsiders will increase prices too, unambiguously resulting in lower consumer surplus.

[9]A parallel strand has emphasized the role of uncertainty in efficiency gains (see, e.g., Choné and Linnemer, 2008; Amir et al., 2009; Cunha et al., 2014).

[10]Katsoulacos and Ulph (2009) also combine the decision-error framework with deterrence and procedural effects.

[11]In a recent work, Vasu (2019) studies the optimal contract that a seller and a buyer offer to their respective advisors who can identify the value of the synergies arising from the merger. However, the paper focuses on the bargaining process between the merging firms and is thus not related to the problem of merger review, to which our work contributes.

1.2 The model

Environment. Suppose that two firms (1 and 2 without loss of generality) contemplate a merger in a N-firm Cournot industry, but in order to do so they need approval by an AA. Before the merger, all firms produce at a (constant) marginal cost $c_0 > 0$. Conditional on the merger being rejected, let π_0 and w_0 denote each firm's profit and the consumer surplus (hereafter, CS), respectively. Similarly, once the merger takes place, depending on the (uncertain) post-merger marginal cost $c \leq c_0$ (more below), the merged firm obtains a profit $\pi_M(c)$ and the consumer surplus is $w_M(c)$. Define the change in the merged entity profits and in the consumer surplus as $\pi(c) \triangleq \pi_M(c) - 2\pi_0$ and $w(c) \triangleq w_M(c) - w_0$, respectively. The merger is said to be profit-increasing if $\pi(c) \geq 0$ and CS-increasing if $w(c) \geq 0$. To make the problem interesting for our purposes, in line with Besanko and Spulber (1993), assume that $\mathrm{E}[w(c)] < 0$ — i.e., behind the veil of ignorance the AA always rejects the merger.

Information acquisition. The (constant) marginal cost c of the merged entity is uncertain, and is distributed uniformly on the support $[0, c_0]$. The difference $c_0 - c \geq 0$ captures the efficiency improvement (cost synergies) produced by the merger. The AA is active (Sørgard, 2009), meaning that the AA can carry out an internal investigation in order to learn the realization of c.[12] The investigation technology is *all-or-nothing*: it is successful with probability $p \in [0, 1]$, while nothing is learned with probability $1 - p$.[13] Before merging, the two firms have no skills to learn the state of nature c, but can delegate an expert (advisor) in order to gather information on their behalf. Conditional on paying the information acquisition cost, the expert learns c with probability p, and remains uninformed otherwise. For simplicity, we assume that the expert cannot falsify information *vis-à-vis* the AA. Hence, he can either disclose the true state of nature if he knows it, or produce an uninformative report.[14] The decision of whether gathering information or not is private information of the expert: a moral hazard problem.
Gathering information is costly both for the AA and the expert, whose (sunk) cost of acquiring information is (for simplicity) identical and equal to $\psi \geq 0$ (see, e.g., Gromb and Martimort, 2007). We assume on purpose that the expert and the AA own the same information acquisition technology in order to isolate results from the direction of such asymmetry (we discuss the role of asymmetries in the information acquisition technologies in Section 1.3.5).

[12] Evidence of active merger policies is discussed in Sørgard (2009).
[13] In Sørgard (2009) the parameter p is interpreted as the probability with which a merger is investigated.
[14] A different information acquisition technology for the expert is discussed in Section 1.3.5.

This information acquisition process is meant to capture the idea that the merger review process, in most jurisdictions, is articulated in two phases. In the first phase the AA carries out a quick review to examine whether the deal entails any potentially serious issues. If this is not the case, the merger is cleared; otherwise, the process enters into the second phase in which the AA requires the proponent firms to produce evidence on the merger efficiency improvements. During this phase, the AA's lawyers and economists conduct internal analysis, request and review documents from parties and non-parties, and interviews take place with interested entities (such as customers, competitors, suppliers, etc.), with the merging parties, and sometimes with business people, the counsel and outside economists.[15]

Contracts. The expert's contract consists of a two-part tariff: a fixed payment $F \geq 0$ plus a linear component $\alpha \geq 0$ contingent on the deal completion (i.e., the merger being approved). This remuneration structure is commonly observed in the real world (see, e.g., Rau, 2000). All our results are unchanged if the contingent fee paid to the expert is specified as a fraction of the change in the merged firm's profit — i.e., in case of merger approval, the expert's remuneration is $F + \beta\pi(c)$, with $\beta \in (0, 1)$. Details are available upon request.

Disclosure regimes and timing. We consider two alternative contract disclosure regimes: one in which the expert's contract must be disclosed to the AA ($j = D$), the other in which the contract is kept secret ($j = S$). The timing of the game is as follows:

- $t = 1$ The state of nature c realizes.
- $t = 2$ Firms 1 and 2 decide whether to propose the merger and (if so) they also decide whether to hire an expert and which contract to offer him. The contract is observed by the AA if its disclosure is mandatory.
- $t = 3$ If hired, the expert decides whether to gather information or not, and so does the AA.
- $t = 4$ The expert submits a report to the AA, who learns the outcome of the internal investigation, updates beliefs and approves or rejects the merger.
- $t = 5$ If the merger is approved, the post-merger outcome realizes (more below), and the fees are paid to the expert; otherwise, the status quo remains.

[15] See, for example, the merger investigation process conducted at the Federal Trade Commission (FTC) (www.ftc.gov/tips-advice/competition-guidance/guide-antitrust-laws/mergers/premerger-notification-merger-review).

Hence, the AA exerts an activity level (information acquisition) with the aim of discovering the merger's welfare effect (see, e.g, Sørgard, 2009). As mentioned above, AAs typically involve market insiders — i.e., competitors and third parties — during the merger review process, to gather information on the market or get comments on the draft of the Statement of Objections (Giebe and Lee, 2019).[16]

More generally, the timing of the model captures the simple idea that merger decisions may take quite a long time, and that discovering and documenting efficiency gains might be costly both for the proponent firms and the AA.[17] In contrast to Nocke and Whinston (2010, 2013) and Sørgard (2009), in our model the AA does not commit to a merger policy, but makes a decision on the basis of the information elicited throughout the review process. This seems to reflect real world practices where a decision is made either right away (i.e., in Phase 1 of the review process) or in Phase 2 after evaluation of the merger (see Section 1.3.5 for a discussion of this issue).

Payoffs. The expert is risk-neutral, but is protected by limited liability — i.e., the fixed fee F cannot be negative — and (without loss of generality) his outside option is normalized to zero. Hence, the expert's payoff is

$$u_e(\cdot) \triangleq \max\{0, F + \alpha d(c) - \psi e\},$$

where, for any state c, $d(c) \in \{0,1\}$ is an index that takes value 1 if the merger is approved and 0 otherwise. Of course, $d(c) = 0$ when the AA's decision is made with no information available about the state of nature since the merger is not CS-increasing in expected terms. Similarly, let $e \in \{0,1\}$ be an index that takes value 1 if the expert gathers information and 0 otherwise.

The joint payoff of the proponent firms is

$$u_f(\cdot) \triangleq (\pi(c) - \alpha)d(c) - F,$$

[16] Giebe and Lee (2019) show that in the EU the number of cases where competitors were given the opportunity to voice their opinions has sharply increased between 1990 and 2013. In the U.S., competitors' claims were traditionally treated restrictively, but both the Department of Justice and the Federal Trade Commission have recently started to widen the extent of competitor participation in merger proceedings by conducting an 'open door' policy.

[17] A study conducted in 2017 by Allen & Overy, an international law company, shows that the cases that are cleared unconditionally in the first phase obtain clearance within 30 working days. By contrast, periods to receive unconditional clearances following an in-depth investigation range from an average of 58 working days (Turkey) to 144 working days (the Netherlands), while prohibitions generally took over 100 working days, with the in-depth investigation period for deals blocked in the U.S. averaging 467 working days. In Europe, the European Commission makes regular use of its *stop the clock powers*, suspending the review by asking the parties for additional information. In the U.S., the average duration of in-depth investigations is consistently increasing over recent years, thus leading the DOJ (Department of Justice) to announce that it wants to increase the speed and reduce the burden of merger reviews.

where, of course, $\alpha = F = 0$ if the expert is not hired.

The AA's preferences are determined by a loss function which takes into account, together with the cost of acquiring information, the probability (computed from ex-ante point of view) of making a wrong decision, weighted by the respective damage to consumers. For any state $c \in [0, c_0]$, let $s(c) \in \{0, 1\}$ be an index that takes value 1 if, in that state, the merger is CS-increasing and 0 otherwise. With the all-or-nothing information structure considered above, for any disclosure regime $j = D, S$, the AA maximizes the following objective function:

$$U^j(\cdot) \triangleq -\int_{c \in \{c: s(c)=1\}} \Pr\left[d(c) = 0 | s(c) = 1, j\right] w(c)\, dc - \psi a.$$

where $a \in \{0, 1\}$ is an index that takes value 1 if the AA acquires information, and 0 otherwise. In words, the AA is only concerned with a type-I error — i.e., the event in which a CS-increasing merger is rejected — since this is the only error that can happen with the *all-or-nothing* information acquisition technology introduced before (we consider type-II errors in Section 1.3.5).

Market structure and technical assumptions. In order to obtain closed form solutions, we consider a (linear) inverse demand function

$$P(\cdot) \triangleq \max\left\{0, A - \sum_{i=1}^{N} q_i\right\},$$

with $c_0 < A$. We argue in Section 1.3.5 that our results hold qualitatively with a more general demand function.

Firm i's production is denoted by q_i. Hence, in the *pre-merger equilibrium*, each firm produces $q_0 \triangleq \frac{A-c_0}{N+1}$. Accordingly, each firm's profit is

$$\pi_0 \triangleq \left(\frac{A - c_0}{N+1}\right)^2,$$

and the consumer surplus is

$$w_0 \triangleq \frac{N^2 \pi_0}{2}.$$

Next, consider the scenario in which the merger is approved. The merged entity will compete in the industry with the remaining $N-2$ firms. To simplify the analysis, we assume that once the merger has occurred the cost c becomes common knowledge to its competitors — i.e., a standard Cournot equilibrium is played in this event (in other words, we abstract from considering additional effects driven by the fact that the merged entity owns private information on its cost structure). Accordingly, following the merger, the equilibrium of the market game is such that the merged entity produces

$$q_M(c) \triangleq \frac{A-(N-1)c+(N-2)c_0}{N},$$

while each other competitor produces

$$q_i(c) \triangleq \frac{A-2c_0+c}{N}, \quad \forall i \geq 3.$$

The profit of the merged firm is

$$\pi_M(c) \triangleq \frac{[A+(N-2)c_0-(N-1)c]^2}{N^2}.$$

Finally, the consumer surplus is

$$w_M(c) \triangleq \frac{[(N-1)A-(N-2)c_0-c]^2}{2N^2}.$$

For simplicity and without loss of insights, in what follows we normalize c_0 to 1.
Conditional on the merger being CS-increasing, let

$$w \triangleq \mathrm{E}[w\,(c)\,|s\,(c)=1]$$

denote the expected change in consumer surplus. Similarly, let

$$\pi \triangleq \mathrm{E}[\pi\,(c)\,|s\,(c)=1],$$

denote the expected change in the merged firm's profit. As we will see, the comparison between the two alternative disclosure regimes yields interesting results in our model if information acquisition technology is sufficiently accurate. Specifically, we impose the following parametric restriction.

Assumption 1.1. $p \geq \dfrac{\pi-2w}{\pi-w}$.

In the Appendix we show that this assumption implies that the number of firms in the industry is not too small.
As a standard tie breaking condition, in all pure strategy equilibria of the game we assume that, whenever the expert or the AA are indifferent between acquiring information and remaining uninformed, they prefer the first option.
Finally, since the game is sequential and there is incomplete information, our equilibrium

concept will be PBE.

1.3 Equilibrium analysis

In this section, we characterize the equilibrium of the game with and without disclosure of the expert's contract. Then, in Section 1.3.4, we compare the expected consumer surplus in the two regimes.

1.3.1 Preliminaries

We first state a few preliminary results that will be useful in the subsequent analysis. First, comparing the outcomes with and without the merger, it is easy to show that the merger is profit-increasing — i.e., $\pi(c) \geq 0$ — if and only if $c \leq \tilde{c}$, where

$$\tilde{c} \triangleq \frac{N^2 + N(\sqrt{2}-1) - 2 - A[N(\sqrt{2}-1) - 1]}{N^2 - 1}.$$

Analogously, the merger is CS-increasing — i.e., $w(c) \geq 0$ — if and only if $c \leq \hat{c}$, where

$$\hat{c} \triangleq \frac{N + 2 - A}{N + 1}.$$

Notice that $\hat{c} < \tilde{c}$, so that if the merger is CS-increasing, it is also profit-increasing, but not the other way around.[18] Hence, without loss of generality, in the following we assume that the merger is always proposed. The reason is that, whenever the AA approves the merger because it is CS-increasing, then it must also be profitable for the proponent firms.[19]
Of course, in order to make the problem interesting, we focus on the region of parameters where $\hat{c} > 0$. In the Appendix, we show that together with $\mathrm{E}[w(c)] < 0$, these inequalities are satisfied for $N \in [A-2, 2(A-2)]$, which we will assume throughout.
Next, since $s(c) = 1$ if and only if $c \leq \hat{c}$, the following results hold true.

Lemma 1.1. *π is decreasing in N, while $\hat{c}$, w and $\hat{c}\pi$ are increasing in N. Moreover, $\pi > 4w$.*

The intuition is straightforward. Clearly, as the number of firms in the market grows larger, the merger becomes less profitable because, once merged, the proponent firms will

[18] Such a result holds true in much more general Cournot models (see, e.g., Nocke and Whinston, 2010).

[19] Clearly, the contingent fee that firms pay to the expert (in equilibrium) is lower than the expected gain π.

be exposed to more competition by the non-merging firms. By contrast, as the market becomes more competitive, it is more likely that the merger is CS-increasing — i.e., $\hat{c}$ being increasing in N — because the market power of the merged firm becomes weaker in regards to the efficiency gain it brings. The same intuition explains why the expected change in consumer surplus w is increasing in N.
We are then ready to characterize the expert's reporting behavior and the AA's decision rule.

Proposition 1.1. *The informed expert sends a truthful report to the AA if and only if $c \leq \hat{c}$, otherwise he produces an uninformative report. The AA clears the merger if and only if the evidence gathered internally or by the expert is such that $c \leq \hat{c}$.*

The intuition is straightforward. Conditional on the expert having learned the state of nature, he has an incentive to produce an informative report if and only if this proves that the merger is CS-increasing because, in this case, the merger will be approved and he will be remunerated. Otherwise, he produces an uninformative report. In this case, however, the AA anticipates the expert's strategic behavior and rejects the proposal unless the evidence collected internally shows that the merger is CS-increasing.

1.3.2 The disclosure regime

When the expert's contract is observable, the game is sequential. Hence, the proponent firms act as a Stackelberg leader to influence the AA's behavior. To save on notation, hereafter we normalize the expert's fixed payment F to zero since he is protected by limited liability.

Information acquisition. We start by analyzing the expert's incentive to gather information. The key point to notice is that because the proponent firms are unable to assess the evidence collected by the expert, they may end up paying him even when he has not gathered information, but the merger was approved only thanks to the evidence collected by the AA. Hence, the expert's incentive to gather information depends on the AA's behavior: other things being equal, the expert may want to save on the cost of acquiring information and let the AA learn the state of nature on his behalf. This creates a standard free-riding problem which will be key to our analysis.
If the AA is expected not to acquire information, the expert has an incentive to gather information if and only if

$$p\hat{c}\alpha - \psi \geq 0 \iff \alpha \geq \alpha_L \triangleq \frac{\psi}{p\hat{c}}.$$

This inequality has a simple interpretation. When only the expert gathers information, the merger is approved if the expert learns the state of nature (with probability p) and if the merger brings strong enough cost synergies — i.e., if $c \leq \hat{c}$. When, instead, neither the expert nor the AA gather information, the merger is always rejected (by assumption) and the expert never gets rewarded.

If the AA is expected to acquire information, the expert gathers information if and only if

$$(1-(1-p)^2)\hat{c}\alpha - \psi \geq p\hat{c}\alpha \iff \alpha \geq \alpha_H \triangleq \frac{\psi}{p(1-p)\hat{c}}.$$

The right-hand side of the inequality is the expert's expected utility if he does not gather information. As mentioned before, conditional on the merger being approved, he can be rewarded even if he has not gathered information. The left-hand side of the inequality is, instead, the expert's expected utility when he gathers information. In this case, what matters is the probability that (conditional on both the expert and the AA gathering information) at least one learns the state of nature — i.e., $1-(1-p)^2$.

It can be easily verified that $\alpha_H > \alpha_L$ and that both these values are decreasing in N. The reason is simple. Clearly, α_H exceeds α_L because the expert has less incentive to acquire information when the AA acquires information, owing to the free-riding problem mentioned before. Moreover, holding the AA's behavior constant, as the market becomes less concentrated, the anticompetitive effect of the merger weakens compared to the efficiency gains it may bring. Hence, the probability of the merger being approved increases so that it is less risky for the expert to incur the information acquisition cost, which in turn weakens his incentive compatibility constraint.

We can then show the following.

Lemma 1.2. *The expert's behavior is as follows:*

(i) If $\alpha \geq \alpha_H$ he gathers information regardless of the AA's behavior.

(ii) If $\alpha < \alpha_L$ he does not gather information regardless of the AA's behavior.

(iii) If $\alpha \in [\alpha_L, \alpha_H)$ the expert gathers information if and only if he expects the AA not to do so.

Since acquiring information is costly, the expert and the AA face a standard free-riding problem. Other things being equal, the expert would like the AA to gather information since this would allow him to save on the cost of getting informed and nevertheless obtain a reward with probability $p\hat{c}$. Yet, if the AA does not collect information, the expert has an incentive to do so on his own as long as the bonus that he receives in case the merger is approved is sufficiently large.

Using the same logic, we now analyze the incentive of the AA to acquire information. Since acquiring information is costly, the AA's objective depends on whether the expert (in the equilibrium) gathers or not information.
First, consider the case in which the expert gathers information. The AA has an incentive to gather information if and only if

$$-(1-p)^2\hat{c}w - \psi \geq -(1-p)\,\hat{c}w \iff \psi \leq \psi_L^a \triangleq p(1-p)\hat{c}w.$$

Notice that $(1-p)^2$ is the probability that neither the expert nor the AA are able to learn the state of nature, while $\hat{c}$ is the probability that the merger is indeed CS-increasing. Hence, $(1-p)^2\hat{c}$ is the probability of rejecting a CS-increasing merger. Clearly, the AA acquires information if the cost of doing so is not too large. Notice, however, that this incentive is stronger when the expected gain in consumer surplus from the merger is higher — i.e., when *ceteris paribus* $\hat{c}w$ increases, hence, from Lemma 1.1, when the market is less concentrated.
Second, consider the case in which the expert does not gather information. The AA has an incentive to gather information if and only if

$$-(1-p)\hat{c}w - \psi \geq -\hat{c}w \iff \psi \leq \psi_H^a \triangleq p\hat{c}w.$$

Once again, the AA would like to be informed if the cost of doing so is (relatively) small and if the expected gain in consumer surplus is large enough. However, when the expert is uninformed, only the probability p that the AA learns the state of nature matters.
It is immediate to see that $\psi_L^a < \psi_H^a$, which again reflects the free-riding problem between the AA and the expert — i.e., the AA has less incentive to get informed when it expects the expert to gather information. We can thus show the following result.

Lemma 1.3. *The AA's information acquisition behavior is as follows:*

(i) If $\psi \leq \psi_L^a$ the AA gathers information regardless of the expert's behavior.

(ii) If $\psi > \psi_H^a$ the AA does not gather information regardless of the expert's behavior.

(iii) If $\psi \in (\psi_L^a, \psi_H^a]$ the AA gathers information if and only if the expert does not.

Clearly, if the cost of acquiring information is sufficiently small, the AA always tries to get informed regardless of the expert's behavior; by contrast, if this cost is too high, the

AA prefers to remain uninformed no matter what the expert does. The most interesting case is the region of parameters in which ψ takes intermediate values — i.e. where the free riding problem bites. In this case, the AA would like to invest in information acquisition, but only if the expert is uninformed.

Advisory contract, hiring decision and equilibrium. We can now turn to characterize the contract offered by the firms to the expert, and then the equilibrium of the game. Since we normalized the expert's outside option to zero, it follows that any bonus $\alpha \geq 0$ is accepted by the expert. Hence, the optimal α satisfies the expert's incentive compatibility constraint as equality — i.e., it makes him indifferent between gathering information and remaining uninformed (for given AA's information acquisition choice). Accordingly, in a pure strategy equilibrium, $\alpha = \alpha_H$, if (in the equilibrium) the AA is expected to acquire information, and $\alpha = \alpha_L$ otherwise.

But, do firms have an incentive to rely on the expert? The answer depends on the AA's information acquisition behavior. If the AA remains uninformed (i.e., for $\psi > \psi_H^a$), the expected profit of the proponent firms is $p\hat{c}(\pi - \alpha_L)$ when the expert gathers information.[20] Hence, he will be hired if and only if

$$p\hat{c}(\pi - \alpha_L) \geq 0 \iff \psi \leq \psi_H^e \triangleq p\hat{c}\pi.$$

By contrast, if the AA gathers information (i.e., for $\psi \leq \psi_L^a$), the expected profit of the proponent firms is $(1-(1-p)^2)\hat{c}(\pi - \alpha_H)$ if they hire the expert, and $p\hat{c}\pi$ otherwise. Hence, the expert is hired if and only if

$$(1-(1-p)^2)\hat{c}(\pi - \alpha_H) \geq p\hat{c}\pi \iff \psi \leq \psi_L^e \triangleq \frac{p(1-p)^2}{2-p}\hat{c}\pi < \psi_H^e.$$

Notice that $\psi_L^e < \psi_H^e$ because when the AA does not gather information the merger is rejected with probability 1 if the expert is not hired.

Finally, consider $\psi \in (\psi_L^a, \psi_H^a]$. Within this region of parameters, from Lemma 1.3, we know that the AA is willing to gather information if and only if the expert does not. Moreover, recall that, from Lemma 1.2, for any contingent fee $\alpha \in [\alpha_L, \alpha_H)$, the expert is willing to gather information if and only if the AA does not. Thus, in these cases (i.e., $\psi \in (\psi_L^a, \psi_H^a]$ and $\alpha \in [\alpha_L, \alpha_H)$), there are two pure strategy equilibria in the information acquisition game: one in which only the expert gathers information, the other in which only the AA gathers information. In addition to these equilibria, there exists a mixed strategy equilibrium, in which both players randomize in their information acquisition

[20] In a pure strategy equilibrium, there is no reason, of course, for hiring the expert when he will not be induced to gather information.

behavior. Clearly, if the firms expect the expert and the AA to coordinate on a pure strategy equilibrium, then, for every $\psi \in (\psi_L^a, \psi_H^a]$, they find it optimal not to hire the expert ($\alpha = 0$), thus letting the AA to bear the information acquisition cost. If, by contrast, they expect the mixed strategy equilibrium to be played in the information acquisition game, they can find it optimal to hire the expert, offering him a contract specifying a contingent fee $\alpha^D \in [\alpha_L, \alpha_H)$ which maximizes their expected profit (given the mixed strategies played by the expert and the AA in the information acquisition stage). In the Appendix we prove that, under Assumption 1.1, whenever the merging firms expect the expert and the AA to coordinate on the mixed strategy equilibrium in the information acquisition game, they find it optimal to hire him, offering a contract $\alpha^D \in [\alpha_L, \alpha_H)$, if and only if $\psi \in (\psi_L^a, \overline{\psi})$, where

$$\overline{\psi} \triangleq p(1 - \sqrt{p})\hat{c}\sqrt{w\pi} \in (\psi_L^a, \psi_H^a).$$

By contrast, for $\psi \in (\overline{\psi}, \psi_H^a]$, regardless of the equilibrium they expect to be played in the information acquisition game, the merging parties find it optimal not to hire the expert. Moreover, in the Appendix we prove that, under Assumption 1.1,

$$\psi_L^e < \psi_L^a < \psi_H^a < \psi_H^e.$$

The following result characterizes the equilibrium of the game with disclosure, confining our attention to the most interesting case in which, when $\psi \in (\psi_L^a, \psi_H^a]$ and $\alpha \in [\alpha_L, \alpha_H)$, mixed strategies are played in equilibrium in the information acquisition game.[21]

Proposition 1.2. *With disclosure of the expert's contract, the equilibrium of the game has the following features.*

(*i*) *The proponent firms hire the expert if $\psi \le \psi_L^e$, $\psi \in (\psi_L^a, \overline{\psi})$ or $\psi \in (\psi_H^a, \psi_H^e]$. Moreover:*

- *If $\psi \le \psi_L^e$ the contingent fee is equal to α_H and both the expert and the AA gather information;*
- *If $\psi \in (\psi_L^a, \overline{\psi})$, the contingent fee is equal to*

$$\alpha^D \triangleq \frac{\psi}{p\hat{c}}\sqrt{\frac{\pi}{w}},$$

[21] The reasons for our interest in the mixed strategy equilibrium will be pointed out in Section 1.3.4.

the expert gathers information with probability

$$\sigma_E^D \triangleq \frac{1}{p} - \frac{\psi}{p^2 \hat{c} w},$$

whereas the AA gathers information with probability

$$\sigma_A^D \triangleq \frac{1}{p} - \frac{\psi}{p^2 \hat{c} \alpha^D} = \frac{1}{p}\left(1 - \sqrt{\frac{w}{\pi}}\right);$$

– *If* $\psi \in (\psi_H^a, \psi_H^e]$ *the contingent fee is equal to* α_L *and only the expert gathers information.*

(*ii*) *The proponent firms do not hire the expert if* $\psi \in (\psi_L^e, \psi_L^a]$, $\psi \in (\overline{\psi}, \psi_H^a)$ *or* $\psi > \psi_H^e$. *Moreover:*

– *If* $\psi \in (\psi_L^e, \psi_L^a]$ *or* $\psi \in (\overline{\psi}, \psi_H^a)$, *the AA gathers information;*

– *If* $\psi > \psi_H^e$ *the AA remains uninformed.*

With disclosure of the expert's contract, the game is *de facto* sequential with the proponent firms acting as a Stackelberg leader *vis-à-vis* the AA. In particular, in the region of parameters where the cost of gathering information is large enough to create a free-riding problem but not too large to make information acquisition unprofitable (i.e., for $\psi \in (\overline{\psi}, \psi_H^a)$), the commitment not to hire the expert forces the AA to gather information, which resolves the free-riding problem in favor of the proponent firms. By contrast, for lower values of the information acquisition cost (i.e., for $\psi \in (\psi_L^a, \overline{\psi})$), both the expert and (for any given contingent fee) the AA gather information with relatively high probability, which implies that, if the firms expect the mixed strategy equilibrium to be played in the following stage of the game, then they prefer to hire the expert so to induce him and the AA to randomize in their information acquisition behavior.
The mixed strategy equilibrium has an interesting comparative statics, which we summarize in the following result.

Proposition 1.3. *The impact of* N *on the mixed strategy equilibrium characterized above is as follows:*

- *The bonus* α^D *is decreasing in* N*;*
- *The probability* σ_A^D *is decreasing in* N, *whereas* σ_E^D *is increasing in* N.

The intuition is rather simple. When the market becomes less concentrated, both the probability of the merger being approved and the expected loss of consumer surplus in absolute value increase. As a consequence, for a given value of the contingent fee, in the information acquisition game, both the expert and the AA acquire information with higher probability. This leads the hiring firms to find it optimal to lower the contingent fee, which in turn has a negative effect on the probability with which the AA gathers information (recall that the AA randomizes in order to make the expert indifferent between gathering information or not). As a result, σ_A^D turns out to be decreasing in N as well.

1.3.3 The no-disclosure regime

When the expert's contract is not disclosed to the AA, the game is simultaneous: the proponent firms choose whether or not hiring the expert and the AA decides whether or not starting an internal investigation, simultaneously. Hence, when the free-riding problem arises, the AA and the proponent firms may fail to coordinate on an equilibrium in which at least one of them gathers information. This potential coordination failure generates scope for a mixed strategy equilibrium, which will be at the heart of the welfare analysis developed in the next section.

Proposition 1.4. *Without disclosure of the expert's contract, if $\psi \in (\psi_L^a, \psi_H^a)$ there are two pure strategy equilibria: one in which only the expert gathers information (and is offered α_L), the other in which the AA gathers information and the expert is not hired ($\alpha = 0$). Moreover, there is also a mixed strategy equilibrium in which the expert is hired with some probability while the AA randomizes between gathering information and remaining uninformed. Specifically:*

- *When the expert is hired he is offered a bonus*

$$\alpha^S \triangleq \frac{(1-p)\psi + \sqrt{4p^2\psi\hat{c}\pi + (1-p)^2\psi^2}}{2p\hat{c}},$$

  which induces him to gather information.

- *The AA collects information with probability*

$$\sigma_A^S \triangleq \frac{\pi - \alpha^N}{\pi - (1-p)(\pi - \alpha^N)}.$$

- *The expert is hired with probability*

$$\sigma_F^S \triangleq \frac{1}{p} - \frac{\psi}{p^2 \hat{c} w}.$$

In all other cases — i.e., for any $\psi \leq \psi_L^a$ and $\psi \geq \psi_H^a$ — the equilibrium is as in the game with disclosure.

When the expert's contract is not disclosed, the AA has to form a conjecture on whether the expert has been hired or not, while the proponent firms must conjecture the AA's information gathering decision. This uncertainty leads to the emergence of multiple equilibria in the region of parameters where the AA and the expert have an incentive to free-ride on one another — i.e., for $\psi \in (\psi_L^a, \psi_H^a)$. In this region of parameters the game without disclosure also features an equilibrium in mixed strategies in which the AA and the proponent firms randomize.
The comparative statics of the mixed strategy equilibrium is shown in the following result.

Proposition 1.5. *The impact of N on the mixed strategy equilibrium characterized above is as follows:*

- *The bonus α^S is decreasing in N;*
- *The probabilities σ_A^S and σ_F^S are increasing in N.*

As discussed before, when the market becomes less concentrated, the anticompetitive effect of the merger weakens. This increases the probability of the merger being approved, whereby making the expert not only more willing to acquire information, but also to do so at a (relatively) lower bonus. Hence, α^N is decreasing in N. For the same reason, when the number of firm is higher, σ_A^S has to increase in order to counterbalance the higher incentive of the expert to gather information. By the same token, other things being equal, as N grows larger the merger is more likely to be CS-increasing, which increases the AA's incentive to acquire information. A higher σ_F^S counterbalances this stronger incentive to acquire information because it exacerbates the AA's incentive to free-ride on the proponent firms. This explains why also σ_F^S is increasing in N.

1.3.4 Optimal disclosure rule

In this section, we characterize the optimal disclosure rule. In doing so, following most of the literature (see, e.g., Whinston, 2007, Nocke and Whinston, 2013), we adopt a consumer surplus standard. Accordingly, we compare across the relevant parameter regions the expected loss in consumer surplus in the equilibria of the two games with and without disclosure.

The first important point to notice is that disclosure plays no role in the region of parameters where the free riding problem does not arise. Obviously, disclosure is neutral to consumer surplus when neither the AA nor the proponent firms acquire information — i.e., for $\psi > \psi_H^e$. In this case, no information is acquired in the equilibrium of both games, and the merger is always rejected. Similarly, disclosure has no effect on consumer surplus in the region of parameters where both the expert and the AA acquire information regardless of what the other does — i.e., for $\psi \leq \psi_L^e$. In this case, the merger is approved and rejected with the same probabilities in both games. The same conclusion applies to $\psi \in (\psi_L^e, \psi_L^a]$ and $\psi \in (\psi_H^a, \psi_H^e]$.

We can thus argue that the most interesting region of parameters is that in which the free-riding problem arises — i.e., when each player remains uninformed when it expects the other to acquire information, and vice versa, namely $\psi \in (\psi_L^a, \psi_H^a)$.

The existence of multiple equilibria, in both disclosure regimes, within this range of the parameters, brings out a selection issue. While there are good reasons to focus on the mixed strategy equilibrium, it also worth remarking that this multiplicity does not reduce the generality of our results. The reason is that, if one of the pure strategy equilibria were played, then the games with and without disclosure would yield the same results, meaning that our welfare conclusions, and thus the policy implications, are valid independently of which equilibrium is played.[22]

We now explain our interest in the mixed strategy equilibria. First, from a game theoretic point of view, the reason to select the mixed strategy equilibria is twofold. A standard argument builds on the so called 'purification' argument (Harsanyi, 1973). According to this logic, mixed strategy equilibria are explained as being the limit of pure strategy equilibria for a perturbed game of incomplete information in which the payoff of every player is his own private information.[23] Alternatively, according to evolutionary game theory, mixed strategy equilibria can be selected based on the idea of a large population

[22]Of course, if one considers only pure strategy equilibria the disclosure of the expert's contract has no effect on consumers. In fact, notice that an equilibrium in which only the AA acquires information is equivalent, from the consumers' point of view, to an equilibrium in which only the expert acquires information since both players have the same information acquisition technology. Hence, regardless of which player acquires information in equilibrium, the expected loss for consumers is the same with and without disclosure of the expert's contract.

[23]The mixed nature of the strategy can be seen as the result of each player playing a pure strategy with threshold values that depend on the ex-ante distribution over the continuum of payoffs that a player can have. For example, one could think of our mixed strategy equilibrium as being the equilibrium of a game where the cost of acquiring information (ψ) is uncertain and private information of each player — i.e., the AA and the expert privately observe their (random) costs of acquiring information.

of players (see, e.g., Fudenberg and Tirole, 1991, Ch. 1).[24] Second, from a competition policy perspective, the mixed strategy equilibria are the most interesting ones because, in both disclosure regimes, the probability of rejecting a CS-increasing merger depends on the market structure (i.e., the number of firms in the industry) through the players' indifference conditions (which pin down the probabilities with which players gather information in equilibrium).

In other words, if, in the more interesting region of parameters, players coordinate on the mixed strategy equilibria, the probability with which the state of nature is learned is endogenous and depends on the market structure. This dependency yields interesting results as we are going to show.

First consider $\psi \in (\psi_L^a, \overline{\psi}]$. In this region of parameters, in both disclosure regimes, mixed strategies are played in equilibrium. Notice that the probability with which the expert gathers information is the same in the two regimes. Hence, the disclosure regime which maximizes consumer surplus is the one in which the AA gathers information with a higher probability.

Next consider $\psi \in (\overline{\psi}, \psi_H^a)$. In this region of parameters, if contracts' disclosure is mandatory, the expert is not hired and the AA gathers information (with probability one). Accordingly, the loss in consumer surplus under disclosure is

$$-\underbrace{(1-p)\hat{c}}_{\Pr[d(c)=0,s(c)=1,D]} w.$$

By contrast, in the no disclosure regime, the mixed strategy equilibrium is played, and the expected loss in consumer surplus is

$$-\underbrace{\hat{c}[\sigma_A^S\sigma_F^S(1-p)^2+(1-p)\left(\sigma_A^S(1-\sigma_F^S)+(1-\sigma_A^S)\sigma_F^S\right)+(1-\sigma_A^S)(1-\sigma_F^S)]}_{\Pr[d(c)=0,s(c)=1,S]}w=$$

$$=-\frac{\alpha^S\psi}{p(p\pi+(1-p)\alpha^S)}.$$

Comparing these expressions we have the following result.

Proposition 1.6. *Let $\psi \in (\psi_L^a, \psi_H^a)$. There exists a threshold $\psi^* > \overline{\psi}$ such that the regime in which the expert's contract is disclosed to the AA increases consumer surplus relative to the regime without disclosure if and only if $\psi \geq \psi^*$. The threshold ψ^* is increasing in N.*

[24] This logic postulates that in a large population of players (firms in our case) some always play one strategy (i.e., in our case, do not hire the expert) while the others play the alternative strategy (hire the expert). An equivalent reasoning can be applied to the regime with disclosed contracts.

Thus, for every $\psi \in (\psi_L^a, \overline{\psi}]$, the disclosure of the expert's contract is harmful from a consumer surplus viewpoint. The reason is as follows. In both regimes, as ψ increases, players have weaker incentives to acquire information which, *ceteris paribus*, implies that the probability with which the AA acquires information is decreasing in ψ. However, in the regime in which the expert's contract is disclosed, firms find it optimal to set a contingent fee which is increasing in ψ, and the probability with which the AA gathers information in that regime is in turn increasing in α. It turns out that these two opposite effects of ψ on the probability with which the AA acquires information are of the same magnitude, implying that the considered probability does not depend on ψ under disclosure. As a consequence, within this region of parameters, in which the cost of acquiring information is small, the AA acquires information with a higher probability in the no disclosure regime as compared to the regime in which the expert's contract is disclosed.
For $\psi \in (\overline{\psi}, \psi_H^a]$, when the cost of acquiring information is small (i.e., ψ is still relatively low), the expert and the AA have (other things being equal) a relatively strong incentive to invest in information acquisition. In particular, in the regime without disclosure the two probabilities σ_A^S and σ_F^S are (relatively) high. Hence, consumers prefer the expert's contract not to be disclosed because this regime minimizes the probability that none of them gets informed (which is equivalent to the probability of rejecting a CS-increasing merger) compared to the case of disclosure where only the AA invests in information acquisition. By contrast, when ψ is large enough, the AA and the expert have a relatively low incentive to be informed when they play mixed strategies in the regime without disclosure — i.e., σ_A^S and σ_F^S are small — while with disclosure the AA acquires information with certainty. As a result, consumers prefer the regime with disclosure because this minimizes the risk of rejecting a CS-increasing merger.
Interestingly, the region of parameters in which disclosure harms consumers expands as the market becomes less concentrated. The reason is as follows. When the number of firms in the industry increases, the anticompetitive effect of the merger weakens. Hence, conditional on being informed, the AA rejects the merger with lower probability. Other things being equal, a lower probability of rejection tends to increase consumer surplus more without disclosure than with disclosure because in the former regime both players acquire information (although they do so with probability less than one) while in the latter only the AA acquires information (with certainty).
In order to provide more insights on the effect of N on the optimal disclosure rule, below we provide a few intuitive illustrations of the comparative statics just discussed. For fixed values of A, p and $\psi \in (\overline{\psi}, \psi_H^a]$, in Figure 1.1 we plot the expected CS loss, in absolute value, in the disclosure regime (continuous line) and in the no-disclosure regime (dashed line), showing that there exists a threshold N^* above which consumers prefer the regime

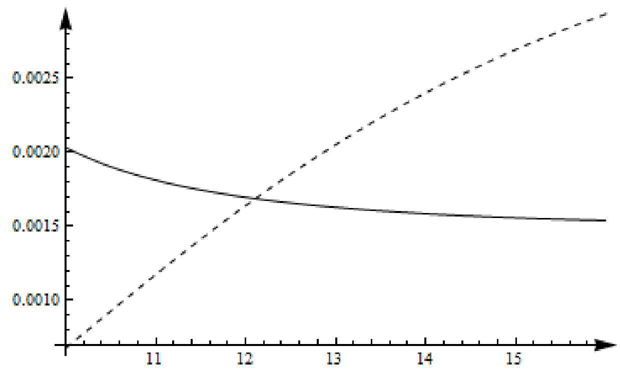

Fig. 1.1: Baseline example. Parameters' values: $A = 10, p = 0.9, \psi = 0.01$. Consumer expected loss (in absolute value) as function of N in the game with disclosure (dashed line) and without disclosure (continuous line).

in which the expert's contract is not disclosed. First notice that in the case of disclosure, the loss in consumer surplus is (in absolute value) increasing with N, while the opposite holds in the regime without disclosure. The reason is simple. When competition increases, the loss in consumer surplus caused by the rejection of a CS-increasing merger increases since the positive effect of the merger on the (equilibrium) price is diluted by a more intense competitive pressure — i.e., $\hat{c}w$ is increasing in N. This effect is present both with and without disclosure. In the latter case, however, there is also the effect of the market structure on the probabilities with which players randomize in the mixed strategy equilibrium. A stronger competitive pressure in this case reduces the expected loss in consumer surplus since it increases the probability with which the proponent firms hire the expert and the AA gathers information. This effect, in fact, outweighs the positive effect of N on $\hat{c}w$ discussed above.

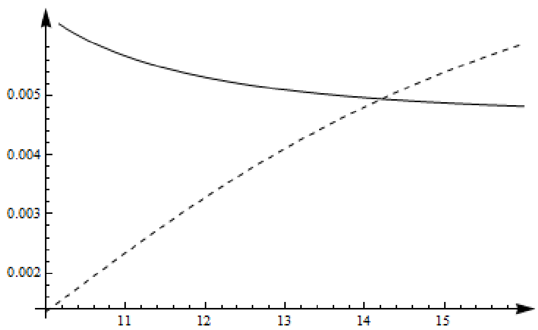

(a) $A = 10, p = 0.9, \psi = 0.02$.

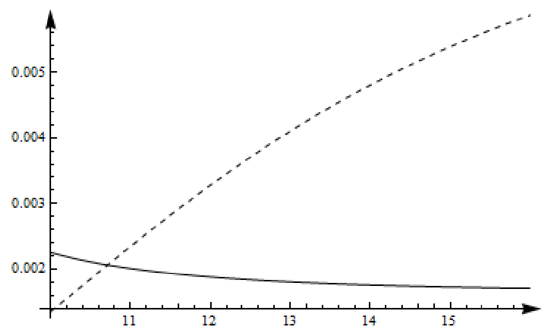

(b) $A = 10, p = 0.8, \psi = 0.01$.

Fig. 1.2: Consumer expected loss (in absolute value) as function of N in the game with disclosure (dashed line) and without disclosure (continuous line), for different values of the parameters.

In Figure 1.2 panel (a), the value assigned to ψ increases and the region in which disclosure harms consumers shrinks. This is because, as information acquisition becomes more costly, the AA and the expert get informed with lower probability in the mixed strategy equilibrium of the regime without disclosure, while the probability of getting informed in the regime with disclosure is constant.

Finally, in Figure 1.2 panel (b), the value assigned to p drops — i.e., the individual probability of learning the state of nature decreases. Again, compared to Figure 1.1, the area in which consumers prefer the regime without disclosure expands. The reason is simple: as p drops, the probability of remaining uninformed in the regime with disclosure increases at a faster rate than in the regime without disclosure in which both players acquire information (although they do so randomly).

1.3.5 Discussion

Asymmetries in the information acquisition technology. One may wonder what happens when the expert and the AA differ in their ability to acquire information. The free-riding problem between the expert and the AA driving our results should not vanish as long as the information acquisition technologies of the two players are not too different. Of course, if acquiring information is costless for one player, in equilibrium this player should acquire information. As intuition suggests, in this case disclosure would have no impact on consumer surplus. Yet, in light of the constraints that antitrust authorities face in real world, it seems implausible to assume that information acquisition is costless for them, and since time is valuable it should be costly for the experts too. In addition, it is hard to imagine why the ability of the AA and the expert should be different in the real world. Indeed, in addition to the fact that merger evaluations are guided by same theory and data that are available to both the AA and the experts, antitrust authorities and consulting companies around the world often recruit people (e.g., graduate students) with the same background and from the same market.

Biased expert and type-II errors. Up until now, we assumed that the AA and the expert use the same information acquisition technology. However, one can argue that the differences in their objectives shape the information acquisition technology employed by the two players. The AA aims to gather unbiased information about the value of the efficiencies, hence it is natural to consider the information acquisition technology described in the baseline model. On the contrary, the expert has always incentives to claim that substantial efficiencies are present — i.e., to produce a report stating that $c \in [0, \hat{c}]$ — in order to persuade the AA to approve the merger. However, the probability with which he is able to produce such a convincing report clearly depends on the state of nature c: the lower the true level of efficiencies, the lower the probability that the expert can find convincing evidence that the merger is CS-increasing. To model this in the simplest possible way, we assume that the outcome of the expert's activity is a binary informative signal $s \in \{0, 1\}$: $s = 0$ means that no evidence of substantial efficiencies is found by the expert, whereas $s = 1$ means that the expert's report credibly claims that the merger is CS-increasing.[25] For simplicity, we posit $\Pr[s = 1|c] \triangleq 1 - c$. To make things interesting, we assume that, upon receiving positive evidence from the expert, and in the absence of any evidence gathered internally, the AA approves the merger.[26] Thus, as in the baseline model, the merger can be approved thanks to the evidence gathered by the AA or by the expert. However, while in the former case the merger clearance decision is taken in a complete information environment, in the latter case, differently from the baseline model, a type-II error may be committed (like, for example, in Besanko and Spulber, 1993, and Sørgard, 2009), since, with probability $\Pr[c > \hat{c}|s = 1] = (1 - \hat{c})^2$, the approved merger is indeed CS-decreasing.[27]

In the Appendix we prove that, under some restrictions on p, the equilibria of the game, under the two disclosure regimes, are as in the baseline model (up to a different specifications for the thresholds on the information acquisition cost and the mixed strategies played in equilibrium). Moreover, by using numerical examples, we show that also the results of Proposition 1.6 apply to this different specification of the model, in which the expert is biased and also type-II errors matter in the AA's choices.

General demand function. One may wonder whether our conclusions hold with a more general downward-sloping demand function. It is well known (Nocke and Whinston, 2010) that in a general Cournot model, under standard assumptions, any merger among $M \geq 2$ firms that the AA would be willing to approve is strictly profitable for the merging parties. Hence, $\hat{c} < \tilde{c}$ even with a more general demand function, so that our main results would hold true qualitatively.

Merger policy. Motivated by the real world practice, we have ruled out the possibility of a merger policy — i.e., an ex-ante commitment of the AA to approve only certain mergers. Besanko and Spulber (1993), Sørgard (2009) and Nocke and Whinston (2010,

[25] Clearly, regardless of the state of nature, no evidence can be produced if the expert does not gather information (i.e., if he does not bear the information acquisition cost ψ).

[26] This condition amounts to assume that the market concentration is relatively low (see the Appendix).

[27] Clearly, as in the baseline model, there is a positive probability of making a type-I error whenever neither the expert nor the AA find evidence, and the merger is not approved. For simplicity, we suppose that the AA equally weighs the expected costs (in terms of loss of consumer surplus) of making a type-I and a type-II error in the merger clearance decision.

2013) argue that an AA acting in the consumers' interest may benefit from committing ex-ante to a clearance rule or to an activity level that is different from the ex-post optimal rule. This result is likely to hold also in our framework since any commitment ex-ante of the AA would modify not only its own incentive to gather information ex-post but also that of the proponent firms. In this sense the AA could exploit the commitment option as a way to exert a first mover advantage that in our model with disclosure is exerted by the proponent firms. Commitment ex-ante to a merger policy that is ex-post inefficient may however be subject to renegotiations issues that should be considered carefully in the analysis. The same 'time inconsistency' issue would emerge if the AA could commit (before the review process) not to gather information in order to solve the free riding problem with the proponent firms. This is clearly outside the scope of the present paper, we hope to address these questions in future research.

Total welfare standard. Consistently with the antitrust law of several countries, and in line with the available literature (e.g., Nocke and Whinston, 2010, 2013), we assumed that the AA adopts a consumer surplus standard. If, on the contrary, we suppose that the AA follows a total welfare standard, namely its objective is the sum of consumer surplus and aggregate profits of the merged entity and the outsider firms, then, under complete information, it would approve the merger if and only if $c \leq c^*$. However, since, in our framework, $c^* > \tilde{c}$, it follows that, whenever the merging firms find it optimal to propose a merger, the AA's optimal clearance decision is to approve it. Therefore, the AA would have never incentives to run costly internal investigations about the efficiencies, and the free riding problem which is key to our analysis would wipe out. Clearly, by continuity, we expect our results to hold when the AA's objective is a weighted sum of consumer surplus and aggregate profits, and the relative weight put on the aggregate profits is relatively small.

Alternative timing. Up until now, we assumed that, once the expert is hired, he decides whether or not to disclose his report to the AA. One may argue that, once the expert finalizes his activity, the report is in the hands of the hiring firms, who then interact with the AA. Would our results change if we assume that, upon receiving a report from the expert, the merging firms decide whether or not to disclose it to the AA? Clearly, our analysis is unchanged if, exactly as the expert would do, the merging firms find it optimal to disclose the expert's report if and only if it states that $c \leq \hat{c}$. This would be the case if firms' and expert's incentives are perfectly aligned, as it happens if the contingent fee is proportional to profits. By the same token, it can be easily verified that, in the extension with biased expert discussed above, the firms always find it optimal to disclose the expert's report when it claims that the merger is CS-increasing, exactly as the expert himself would do.[28] By contrast, in the baseline model it is not obvious that firms always find it optimal to disclose any report claiming that $c < \hat{c}$. The reason is as follows. Without imposing further parametric restrictions, it may be that $\pi(\hat{c}) < \pi$ is lower than the contingent fee to be paid in equilibrium. Therefore, once the expert learns that c is slightly lower than $\hat{c}$, the firms may find it optimal to not to propose the merger. However, in these cases, whenever the merging firms find it optimal to propose a merger, the AA's optimal clearance decision is to approve it, thus making our analysis uninteresting.

Role of the advisor. To understand the role played by the expert in our model, notice that, in a benchmark case without an advisor, in which the AA's final decision is based only on its own investigation process, the agency would gather information for every $\psi \leq \psi_H^a$. Hence, from a consumer surplus viewpoint, in both disclosure regimes, the possibility of hiring an expert to certify efficiencies unambiguously benefits consumers for $\psi \leq \psi_L^e$ and $\psi \in [\psi_H^a, \psi_H^e]$. By contrast, for $\psi \in (\psi_L^a, \psi_H^a)$, the presence of the expert may harm consumers. To see this, consider the no disclosure regime: from the above analysis, it follows that, for $\psi > \psi^*$, consumers would be better off in the absence of an expert, since in that case there would not be a free-riding problem in the information acquisition process.

Passive AA. We now consider the polar opposite case, in which the AA has a passive role — i.e., it takes its merger clearance decision based only on the evidence provided by the expert. In this case, the merging firms would hire the expert, with a contingent fee α_L to induce him to gather information, for every $\psi \leq \psi_H^e$. Hence, from a consumer surplus viewpoint, in both disclosure regimes, giving the AA the possibility of undertaking internal investigations on the merger proposal unambiguously benefits consumers for $\psi \leq \psi_L^e$. By contrast, for $\psi \in (\psi_L^a, \psi_H^a)$, an active role played by the AA in the information acquisition process may harm consumers. To see this, consider the no disclosure regime: from the above analysis, it follows that, for $\psi > \psi^*$, consumers would be better off in the presence of a passive agency. Thus, in this region of parameters, the AA's ability to commit (before the review process) not to gather information would be beneficial to consumers.

[28] In fact, it can be easily proved (see the Appendix) that $\mathrm{E}[\pi(c)|s = 1] > \mathrm{E}[w(c)|s = 1] > 0$ and that, in the equilibrium, the fee to be paid to the expert in case of merger approval is lower than $\hat{\pi}$. Moreover, firms are indifferent between disclosing or not an uninformative report ($s = 0$), since, in each case, the AA's approval decision would not be affected by this choice. Hence, in a PBE of the game with this modified timing, firms disclose the expert's report if and only if it claims that the merger is CS-increasing ($s = 1$) and the AA clears the merger if either $s = 1$ and no information is gathered internally or, regardless of the expert's report, it learns that $c \leq \hat{c}$.

Model applicability. To conclude, we point out that, although throughout the paper we focused on M&A advisory, our analysis applies to other types of regulatory issues in which there are no monetary transfers between the regulator and the firms. Specifically, we can think of other firms' potentially anti-competitive conducts subject to antitrust scrutiny (e.g., vertical restraints). In fact, in many cases, as in the merger review process, the key point of the scrutiny consists in the assessment of an unknown state of nature, which determines the impact of the scrutinized conduct on consumer welfare (as well as on firms' profits); firms, who bear the burden of proof, hire economic consultants to obtain the AA's approval of their conduct, and the agency itself may run internal investigations.

1.4 Conclusion

We have developed a simple model showing that when two merging firms, who are uninformed on the uncertain post-merger cost efficiencies, have an incentive to hire an external expert who collects information on their behalf, consumers surplus may drop when the antitrust authority in charge of approving or rejecting the merger is informed about the expert's contract. The negative effect of disclosure on consumer surplus hinges on a novel free-riding problem between the expert and the AA in the information acquisition game, and it is more relevant in highly competitive industries. This suggests that mandatory rules forcing firms that wish to merge to disclose if they are advised by external experts (and how much they pay for such advice) should be used in (relatively) more concentrated markets. By contrast, in competitive industries the AA should actually commit not to learn this information. Given the relevance of advice in M&A practices, this result not only brings a novel theoretical twist but it also provides useful guidelines to know how an optimal review process should take into account information not directly linked to the structure of the industry.

Finally, although the model is very stylized and framed in a typical IO framework, its novel implications are likely to hold in more general environments where firms compete in a Cournot fashion. In addition, its structure can be used to model not only mergers but also other types of regulatory issues in which there are not monetary transfers between the authority and the firms.

Appendix

Appendix 1.A. Proofs

Proof of Lemma 1.1. Since c uniformly distributes over $[0,1]$, it is immediate to show that $\hat{c} > 0$ if and only if $N > A - 2$, and that $\mathrm{E}[W(c)] < 0$ if and only if

$$A^2(2N^2 - 1) - 3A(N^3 + 5N^2 - N - 3) + 3N^3 + 8N^2 - 5N - 7 > 0,$$

which is always satisfied for $N \leq 2(A-2)$, for all $A \geq 3$. Moreover, it can be immediately seen that

$$\frac{\partial \hat{c}}{\partial N} = \frac{A-1}{(1+N)^2} > 0.$$

Next, we have

$$\pi = \frac{A^2(N^2 + 4N + 1) + 4A(N^3 - 3N - 1) + N^4 - 4N^3 - 3N^2 + 8N + 4}{3N^2(N+1)^2}$$

and

$$w = \frac{-A^2(3N^2 - 1) + A(3N^3 + 9N^2 - 2N - 4) - (N+2)(3N^2 - N - 2)}{6N^2(N+1)^2}.$$

To obtain the comparative static results, first notice that $\hat{c} \geq 0$ and $\mathrm{E}[W(c)] \leq 0$ for $A \in [\hat{A}, N+2]$, where

$$\hat{A} \triangleq \frac{(N+1)\sqrt{3(3N^4 + 2N^2 - 1)} + 3(N^3 + 5N^2 - N - 3)}{6(2N^2 - 1)}.$$

We then compute

$$\frac{\partial w}{\partial N} = \frac{A^2(6N^3 - 4N - 2) - A(3N^4 + 15N^3 - 6N^2 - 18N - 8) + 3N^4 + 7N^3 - 12N^2 - 20N - 8}{6N^3(1+N)^3},$$

which is positive if

$$A > \frac{3N^4 + 15N^3 - 6N^2 - 18N - 8 + N(N+1)\sqrt{9N^4 + 12N^2 + 24N + 4}}{4(3N^3 - 2N - 1)}.$$

This inequality is indeed implied by $A \geq \hat{A}$. Hence, in the relevant range of the parameters, $\frac{\partial w}{\partial N} > 0$.

We then have

$$\frac{\partial \pi}{\partial N} = -\frac{2A^2(N^3+6N^2+4N+1)+2A(N^4-N^3-9N^2-7N-2)-3N^4-N^3+4(3N^2+3N+1)}{3N^3(N+1)^3},$$

which is negative if

$$A > \frac{N(N+1)\sqrt{N^4-N^3+3N^2+5N+1}-N^4+N^3+9N^2+7N+2}{N^3+6N^2+4N+1}.$$

This inequality is indeed implied by $A \geq \hat{A}$. Hence, in the relevant range of the parameters, $\frac{\partial \pi}{\partial N} < 0$.

Next, we compute

$$\frac{\partial \hat{c}\pi}{\partial N} = \frac{x_3A^3+x_2A^2+x_1A+x_0}{3N^3(1+N)^4},$$

where

$$x_3 \triangleq 3N^3+16N^2+9N+2, \quad x_2 \triangleq 3(2N^4-7N^3-28N^2-17N-4),$$

$$x_1 \triangleq -3(N^5+6N^4-13N^3-48N^2-32N-8), \quad x_0 \triangleq 5N^5+18N^4-17N^3-80N^2-60N-16.$$

We have

$$\frac{\partial^2 \hat{c}\pi}{\partial N \partial A} = \frac{1}{3N^3(1+N)^4}\big(3A^2(3N^3+16N^2+9N+2)$$

$$+6A(2N^4-7N^3-28N^2-17N-4)-3(N^5+6N^4-13N^3-48N^2-32N-8),$$

which is positive for

$$A > \frac{-2N^4+7N^3+28N^2+17N+4+N(N+1)\sqrt{7N^4-8N^3+12N^2+12N+1}}{3N^3+16N^2+9N+2},$$

which is again implied by $A \geq \hat{A}$. Thus, $\frac{\partial \hat{c}\pi}{\partial N}$ is increasing in A. Moreover, it can be easily checked that $\frac{\partial \hat{c}\pi}{\partial N}\big|_{A=\hat{A}} > 0$. Hence, we conclude that, under our parametric restrictions, $\frac{\partial \hat{c}\pi}{\partial N} > 0$.

Finally, we compute:

$$\pi - 4w = \frac{A^2(7N^2+4N-1)-2A(N^3+9N^2+4N-2)+N^4+2N^3+7N^2-4}{3N^2(N+1)^2}.$$

Notice that, since

$$\frac{\partial(\pi-4w)}{\partial A} = \frac{2}{3}\frac{A(7N^2+4N-1)-(N^3+9N^2+4N-2)}{N^2(N+1)^2} > 0 \quad \forall A \geq \hat{A},$$

and

$$(\pi - 4w)|_{A=1} = \frac{N^2 - 2N - 1}{3N^2} > 0 \quad \forall N > 2,$$

it follows that, under our parametric restrictions, $\pi - 4w > 0$. ■

Proof of Proposition 1.1. If the AA learns the state of nature the merger is approved if and only if $c \leq \hat{c}$, regardless of the expert's choice. Next, suppose that the AA has learned nothing. Of course, the merger is approved if the expert reports $c \leq \hat{c}$. Thus, the expert finds it optimal to disclose his information in this case. Clearly, if instead he learns that $c > \hat{c}$, then he does not disclose this information. However, when the expert submits an uninformative report, the merger is rejected. This is because the AA draws a negative inference — i.e., it knows that the expert has either discovered that $c > \hat{c}$ or has not found any hard evidence — and we assumed that behind the veil of ignorance the merger is CS-decreasing in expectation. ■

Proof of Proposition 1.2. From Lemma 1.2 and Lemma 1.3, it follows that the pure strategy Nash equilibria of the static game between the expert and the AA are as follows.

1. If $\alpha < \alpha_L$, then:
   - for $\psi \leq \psi_H^a$, only the AA gathers information;
   - for $\psi > \psi_H^a$, neither the expert nor the AA gather information.
2. If $\alpha_L \leq \alpha < \alpha_H$, then:
   - for $\psi \leq \psi_L^a$, only the AA gathers information;
   - for $\psi_L^a < \psi \leq \psi_H^a$ there are two pure strategy equilibria: one in which only the AA gathers information, the other in which only the expert gathers information. Moreover, there exists a mixed strategy equilibrium, characterized as follows. The expert gathers information with a probability σ_E^D such that the AA is indifferent between acquire information or not — i.e.,

     $$-\sigma_E^D(1-p)^2\hat{c}w - (1-\sigma_E^D)(1-p)\hat{c}w - \psi = -\sigma_E^D(1-p)\hat{c}w - (1-\sigma_A^D)\hat{c}w.$$

     Analogously, for any (disclosed) value of α, the AA gathers information with probability σ_A^D such that the expert is indifferent between acquire information or not — i.e.,

     $$\sigma_A^D(1-(1-p)^2)\hat{c}\alpha + (1-\sigma_A^D)p\hat{c}\alpha - \psi = \sigma_A^D p\hat{c}\alpha + (1-\sigma_A^D)0,$$

whose solution can be denoted by $\sigma_A^D(\alpha)$.

- for $\psi > \psi_H^a$, only the expert gathers information.

3. If $\alpha \geq \alpha_H$, then:

   - for $\psi \leq \psi_L^a$, both the expert and the AA gather information;
   - for $\psi > \psi_L^a$, only the expert gathers information.

Therefore:

- For $\psi \leq \psi_L^a$, the proponent firms find it optimal to hire the expert if $\psi \leq \psi_L^e$.
- For $\psi > \psi_H^a$, the proponent firms find it optimal to hire the expert if $\psi \leq \psi_H^e$.
- For $\psi_L^a < \psi \leq \psi_H^a$, the proponent firms would never find it optimal to hire the expert if they expect the expert and the AA to coordinate on a pure strategy equilibrium in the information acquisition game. If, instead, firms expect the other two players to play the mixed strategy equilibrium, a candidate optimal contingent fee α^D solves[29]

$$\max_{\alpha^D \in [\alpha_L, \alpha_H)} [\sigma_A^D(\alpha)\sigma_E^D(1-(1-p)^2) + \sigma_A^D(\alpha)(1-\sigma_E^D)p + (1-\sigma_A^D(\alpha))\sigma_E^D p]\hat{c}(\pi - \alpha),$$

  whose first-order condition yields[30]

$$\alpha^D = \sqrt{\frac{\pi(1-p\sigma_E^D)\psi}{\hat{c}p}}.$$

  After substituting the value for σ_E^D, it can be easily checked that, since $\pi > w$, $\alpha^D > \alpha_L$ and

$$\alpha^D < \alpha_H \iff p > 1 - \sqrt{\frac{w}{\pi}},$$

  which is implied by Assumption 1.1.
  Clearly, the considered solution is the SPNE of the game if and only if it yields a higher expected profit for the merging firms than the equilibrium in which they do not hire the expert and the AA gathers information. In the former case, firms' profit is

$$\hat{c}\pi + \frac{\psi}{p}\left(\frac{\psi}{p\hat{c}w} - 2\sqrt{\frac{\pi}{w}}\right),$$

[29] Clearly, due to the moral hazard problem in the relationship between the merging firms and the expert, the latter gets paid in this mixed strategy equilibrium even when he does not gather information.

[30] It can be easily verified that the corresponding second-order condition is satisfied.

which is higher than the profit that they would get, in expectation, if the expert is not hired (i.e., $p\hat{c}\pi$) if and only if $\psi < \overline{\psi}$.[31]

Finally, to get the SPNE of the game, we must be able to compare the thresholds on the information acquisition cost. From Lemma 1.1, it follows $\psi_H^a < \psi_H^e$. Moreover, $\psi_L^e < \psi_L^a$ if and only if Assumption 1.1 holds. ■

Proof of Proposition 1.3. It can be immediately seen that α^D is increasing in π and decreasing in $\hat{c}$ and w, hence, from Lemma 1.1, it is decreasing in N. By the same token, σ_A^D is decreasing in N since it is decreasing in w and increasing in π, whereas σ_E^D is increasing in N since it is increasing in $\hat{c}$ and w. ■

Proof of Proposition 1.4. Assume that the AA gathers information with probability $\sigma_A \in [0,1]$ and the firms hire the expert with probability $\sigma_F \in [0,1]$.
It is easy to see that, given the firms' strategy, the AA best reply correspondence is

$$\sigma_A(\sigma_F) = \begin{cases} 1 & \text{if} \quad \sigma_F < \sigma_F^S \\ [0,1] & \text{if} \quad \sigma_F = \sigma_F^S \\ 0 & \text{if} \quad \sigma_F > \sigma_F^S \end{cases}$$

where

$$\sigma_F^S \triangleq \frac{1}{p} - \frac{\psi}{p^2\hat{c}w}$$

is such that $\sigma_F^S \in (0,1)$ for $\psi_L^a < \psi < \psi_H^a$, $\sigma_F^S > 1$ for $\psi < \psi_L^a$, $\sigma_F^S < 0$ for $\psi > \psi_H^a$.
Analogously, the firms' best reply correspondence is

$$\sigma_F(\sigma_A) = \begin{cases} 1 & \text{if} \quad \sigma_A < \sigma_A^S \\ [0,1] & \text{if} \quad \sigma_A = \sigma_A^S \\ 0 & \text{if} \quad \sigma_A > \sigma_A^S \end{cases}$$

where

$$\sigma_A^S \triangleq \frac{\pi - \alpha}{\pi - (1-p)(\pi - \alpha)}$$

is such that $\sigma_A^S > 0$ for all $\alpha < \pi$, $\sigma_A^S < 1$ for $\alpha > \frac{1-p}{2-p}\pi \in [0, \frac{\pi}{2}]$.
Then, α^S is obtained by imposing the expert's incentive compatibility constraint to be

[31] It is easy to see that $\pi > 4w$ implies $\overline{\psi} \in (\psi_L^a, \psi_H^a)$.

binding when $\sigma_A = \sigma_A^S \in (0,1)$ — i.e.,

$$\sigma_A^S(1-(1-p)^2)\hat{c}\alpha + (1-\sigma_A^S)p\hat{c}\alpha - \psi = \sigma_A^S p\hat{c}\alpha + (1-\sigma_A^S)0.$$

We have $\alpha^S < \pi \Leftrightarrow \psi < \psi_H^e$ and $\alpha^S > \frac{1-p}{2-p}\pi \Leftrightarrow \psi > \psi_L^e$. Thus, for $\psi < \psi_L^e$: $\sigma_A^S > 1$, and the value of α which makes the expert's incentive compatibility constraint binding is α_H; for $\psi > \psi_H^e$: $\sigma_A^S < 0$, and the value of α which makes the expert's incentive compatibility constraint binding is α_L.
By combining the two best reply correspondences, we get the Nash Equilibria of the game in the no-disclosure regime. ■

Proof of Proposition 1.5. We have

$$\frac{\partial \frac{\alpha^S}{\pi}}{\partial[\hat{c}\pi]} = -\frac{(2p^2[\hat{c}\pi] + \psi(1-p)^2)\sqrt{\psi(4p^2[\hat{c}\pi] + \psi(1-p)^2)} + \psi(1-p)(4p^2[\hat{c}\pi] + \psi(1-p)^2)}{2p[\hat{c}\pi]^2(4p^2[\hat{c}\pi] + \psi(1-p)^2)} < 0.$$

Thus, from Lemma 1.1, α^S is the product of two positive and decreasing functions of N, hence it is decreasing in N.
After substituting α^S into σ_A^S, we get

$$\sigma_A^S = \frac{\psi(1-p) + 2p^2[\hat{c}\pi] - \sqrt{\psi\left(4p^2[\hat{c}\pi] + (1-p)^2\psi\right)}}{2p^3[\hat{c}\pi]},$$

and we can compute

$$\frac{\partial \sigma_A^S}{\partial[\hat{c}\pi]} = \frac{\psi}{2p^3[\hat{c}\pi]^2}\left(2p^2[\hat{c}\pi] + (1-p)^2\psi\sqrt{4p^2\psi[\hat{c}\pi] + (1-p)^2\psi^2} + p - 1\right) > 0.$$

Thus, from Lemma 1.1, σ_A^S is increasing in N. Finally, from Lemma 1.1, it immediately follows that σ_F^S is increasing in N. ■

Proof of Proposition 1.6. For $\psi \in (\psi_L^a, \overline{\psi})$, as argued in Section 1.3.4, it is sufficient to compare the mixed strategies employed by the AA in the two disclosure regimes. Notice that σ_A^S is a decreasing function of ψ. On the contrary, σ_A^D does not depend on ψ. Therefore, to establish that, in the considered range of the parameters, consumers are always better off when the advisory contract is not disclosed, we only need to prove that $\sigma_A^D < \sigma_A^S|_{\psi=\overline{\psi}}$, where

$$\sigma_A^S|_{\psi=\overline{\psi}} = \frac{2p\sqrt{\pi} + \left(p^{3/2} - p - \sqrt{p} + 1\right)\sqrt{w} - \sqrt{\left(1-\sqrt{p}\right)\left(\left(1-\sqrt{p}\right)^3\left(\sqrt{p}+1\right)^2 w + 4p\sqrt{\pi w}\right)}}{2p^2\sqrt{\pi}}.$$

After some algebra, we find that this is the case if and only if

$$\frac{\left(p^{3/2}+p-\sqrt{p}+1\right)^2}{1-\sqrt{p}} > 4p\sqrt{\pi}+\left(\sqrt{p}+1\right)^2\left(1-\sqrt{p}\right)^3\sqrt{w}.$$

Clearly, since $\pi > w$, a sufficient condition in order for the above inequality to hold is

$$\frac{\left(p^{3/2}+p-\sqrt{p}+1\right)^2}{\left(1-\sqrt{p}\right)\left(\left(\sqrt{p}+1\right)^2\left(1-\sqrt{p}\right)^3+4p\right)} > \sqrt{\pi}.$$

Finally, to see that this latter inequality is satisfied, notice that the supremum of the values that the right-hand side can take is $\pi|_{N\to A-2} = 2$, which is lower than the value taken by the left-hand side at $p = \frac{1}{2} < \frac{\pi-2w}{\pi-w}$.
When $\psi \in (\psi_L^a, \overline{\psi})$, comparing the expected consumer loss in the two regimes, it is easy to obtain that consumers are better off without disclosure of the expert's contract if and only if

$$(1-p)\psi^2 + 2(1-p)p^3\hat{c}^2\pi w - \psi\sqrt{4p^2\psi\hat{c}\pi + (1-p)^2\psi^2} > 0.$$

To establish our result, first notice that the left-hand side of the considered inequality is decreasing in ψ.[32] Next, after simple algebra, it follows that the left-hand side of the above inequality takes a negative value when evaluated at $\psi = \psi_H^a$ for all $p > \frac{2\pi-w+\sqrt{w(w+4\pi)}}{2\pi}$, which is implied by Assumption 1.1. Moreover, the left-hand side of the considered inequality takes a positive value when evaluated at $\psi = \overline{\psi}$ if and only if

$$\frac{p\left(2\left(p+\sqrt{p}\right)-1\right)+1}{1-\sqrt{p}} > \sqrt{\frac{\pi}{w}}.$$

Since the left-hand side of this inequality is increasing in p, and $\pi > w$, it follows that the considered inequality is implied by the following one

$$\frac{p(4p-1)+1}{1-p} > \frac{\pi}{w}.$$

Again, the left-hand side of this inequality is increasing in p. Finally, it can be easily checked that $\pi > 3w$ implies that the above inequality is satisfied at $p = \frac{\pi-2w}{\pi-w}$. This proves the existence a threshold $\psi^* \in (\overline{\psi}, \psi_H^a)$ such that consumers are better off in the no-disclosure regime if and only if $\psi < \psi^*$.
Finally, to see that ψ^* is increasing in N, it is sufficient to notice that, when $\psi \in (\overline{\psi}, \psi_H^a)$, the loss in consumer surplus under disclosure is, from Lemma 1.1, decreasing in N, and

[32] This is the case since, in the region of parameters under consideration, the loss in consumer surplus does not depend on ψ under disclosure, whereas it is decreasing in ψ in the no disclosure regime.

the loss in consumer surplus in the no disclosure regime can be rewritten as

$$-\frac{2\psi^2}{p\left(\sqrt{\psi\left(4p^2\hat{c}\pi+(1-p)^2\psi\right)}+(1-p)\psi\right)}$$

which, from Lemma 1.1, is clearly increasing in N. ■

Appendix 1.B. Biased expert and type-II errors

In this Appendix, we provide a detailed analysis of the extension with biased expert outlined in Section 1.3.5 of the paper.

1.B.1. Analysis and Results

In the baseline model, we assumed that the AA and the expert use the same information acquisition technology. Nevertheless, the two players have different objectives: on the one hand, the expert is paid conditional on the merger being approved, regardless of the efficiency gains it brings; on the other hand, the AA is concerned with the expected loss of consumer surplus following its merger clearance decision. Hence, one can argue that these differences in their objectives shape the information acquisition technology employed by the two players.[33]

As for the AA, which clearly aims to gather unbiased information about the value of the efficiencies, it is natural to consider the same information acquisition technology described in the baseline model — i.e., conditional on gathering information, the AA's signal is fully revealing with probability p and uninformative otherwise.

On the contrary, the expert has always incentives to claim that substantial efficiencies are present — i.e., to produce a report stating that $c \in [0, \hat{c}]$ — in order to persuade the AA to approve the merger. However, the probability with which he is able to produce such a convincing report clearly depends on the state of nature c: the lower the true level of efficiencies, the lower the probability that the expert can find convincing evidence that the merger is CS-increasing. To model this in the simplest possible way, we assume that the outcome of the expert's activity is a binary informative signal $s \in \{0, 1\}$: $s = 0$ means that no evidence of substantial efficiencies is found by the expert, whereas $s = 1$ means that the expert's report credibly claims that the merger is CS-increasing. Clearly, regardless of the state of nature, no evidence can be produced if the expert does not

[33] Nonetheless, to keep the analysis as much tractable as possible, we maintain the assumption that the information acquisition cost is equal for the two players.

gather information (i.e., if he does not bear the information acquisition cost ψ). If, on the contrary, the expert gathers information, the conditional probability of finding evidence that the merger is CS-increasing, given the state of nature $c \in [0, 1]$, is

$$\Pr[s = 1|c] \triangleq 1 - c.$$

Hence, recalling that c is uniformly distributed in $[0, 1]$,

$$\Pr[s = 1] = \int_0^1 \Pr[s = 1|c]dc = \frac{1}{2}.$$

To make things interesting, we assume that, upon receiving positive evidence from the expert, and in the absence of any evidence gathered internally, the AA approves the merger — i.e., $\mathrm{E}[w(c)|s = 1] > 0$. This amounts to impose that the market concentration is relatively low.

Thus, as in the baseline model, the merger can be approved thanks to the evidence gathered by the AA or by the expert. However, while in the former case the merger clearance decision is taken in a complete information environment, in the latter case, differently from the baseline model, a type-II error may be committed (like, for example, in Besanko and Spulber, 1993, and Sørgard, 2009), since, with probability $\Pr[c > \hat{c}|s = 1] = (1 - \hat{c})^2$, the approved merger is indeed CS-decreasing. Finally, as in the baseline model, there is a positive probability of making a type-I error whenever neither the expert nor the AA find evidence, and the merger is not approved (more below).

We are now ready to analyze the information acquisition stage. The expert's information acquisition behavior is as follows. If the AA gathers information, the expert knows that, if he gathers information as well, incurring the cost ψ, the merger will be approved not only when the AA finds that $c \leq \hat{c}$ (i.e., with probability $p\hat{c}$), but also when the AA does not find any evidence, whereas the expert does, which happens with probability $(1 - p)\frac{1}{2}$. Hence, the expert is willing to gather information, when also the AA is expected to do so, if and only if

$$\left(p\hat{c} + (1 - p)\frac{1}{2}\right)\alpha - \psi \geq p\hat{c}\alpha \iff \alpha \geq \alpha_H \triangleq \frac{2\psi}{1 - p}.$$

If, on the contrary, the AA does not gather information, the expert knows that the merger can never be approved if he does not bear the information acquisition cost, whereas, bearing such a cost, the merger is approved when $s = 1$ (i.e., with probability $\frac{1}{2}$). Hence, conditional on the AA not gathering information, the expert finds it optimal to acquire information if and only if

$$\frac{1}{2}\alpha - \psi \geq 0 \iff \alpha \geq \alpha_L \triangleq 2\psi,$$

with $\alpha_L < \alpha_H$.

We now turn to analyze the AA's information acquisition behavior. Clearly, if the expert does not gather information, the AA's strategy is as in the baseline model — i.e., it gathers information if and only if $\psi \leq \psi_H^a \triangleq p\hat{c}w$. Next suppose that the expert gathers information. In this case, the ex-ante expected loss of consumer surplus if the merger clearance decision is based (only) on the expert's report, in absolute value, is

$$l \triangleq \underbrace{\Pr[s=0] \times \Pr[c < \hat{c}|s=0]}_{\text{Ex-ante probability of type-I Error}} \quad \times \quad \underbrace{\int_0^{\hat{c}} w(c)f(c|s=0)dc}_{\text{Conditional expectation of CS loss}} \quad +$$

$$+ \underbrace{\Pr[s=1] \times \Pr[c > \hat{c}|s=1]}_{\text{Ex-ante probability of type-II Error}} \quad \times \quad \underbrace{\int_{\hat{c}}^{1} -w(c)f(c|s=1)dc}_{\text{Conditional expectation of CS loss}} \ ,$$

where $f(c|s)$ is the posterior probability density function of c, conditional on the realization of the expert's signal. By denoting with

$$w_I \triangleq \int_0^{\hat{c}} w(c)f(c|s=0)dc > 0, \qquad w_{II} \triangleq \int_{\hat{c}}^{1} w(c)f(c|s=1)dc < 0$$

the expected losses of consumer surplus, given that a type-I and a type-II error are committed, l can be rewritten as follows:

$$l = \frac{1}{2}(\hat{c}^2 w_I - (1-\hat{c})^2 w_{II}).$$

Hence, conditional on the expert gathering information, the AA is willing to do so if and only if

$$-(1-p)l - \psi \geq -l \iff \psi \leq \psi_L^a \triangleq pl.$$

In brief, by gathering information, the AA incurs the information acquisition cost ψ, but the expected loss of consumer surplus l will be borne only when its investigations are unfruitful (i.e., with probability $1-p$). It can be shown that $\psi_L^a < \psi_H^a$.

We can now analyze the firms' hiring decision. If the merger is approved thanks to the evidence provided by the expert, the firms' expected profit is denoted by $\hat{\pi} \triangleq \mathrm{E}[\pi(c)|s = 1] > 0$, whereas, if the approval decision is based on the evidence gathered by the AA, the firms' expected profit is π, as defined in the baseline model. Therefore, if the AA does not gather information, the firms find it optimal to hire the expert, offering a contingent fee $\alpha = \alpha_L$ (to induce him to gather information), if and only if

$$\frac{1}{2}(\hat{\pi} - \alpha_L) \geq 0 \iff \psi \leq \psi_H^e \triangleq \frac{1}{2}\hat{\pi}.$$

By contrast, if the AA gathers information, the firms find it optimal to hire the expert, offering a contingent fee $\alpha = \alpha_H$ (to induce him to gather information), if and only if

$$p\hat{c}(\pi - \alpha_H) + (1-p)\frac{1}{2}(\hat{\pi} - \alpha_H) \geq p\hat{c}\pi \iff \psi \leq \psi_L^e \triangleq \frac{(1-p)^2\hat{\pi}}{2(1-p(1-2\hat{c}))},$$

with $\psi_L^e < \psi_H^e$. Moreover, it can be proved that $\psi_H^e > \psi_H^a$.
As in the baseline model, to make the comparison between the two disclosure regimes interesting, we shall assume that p is relatively high. Specifically, as in the baseline analysis, we consider values of p such that $\psi_L^a < \psi_L^e < \psi_H^a < \psi_H^e$. It follows that the pure strategy equilibria, in both the disclosure regimes, are as in the baseline model. However, since the players' payoffs are now specified in a different way, clearly the mixed strategy equilibria, which, in both regimes, can be defined for $\psi \in (\psi_L^a, \psi_H^a)$, change as compared to the baseline model.
Specifically, following the same logic of the baseline analysis, it can be proved that, in the regime with disclosure of the expert's contract, it can exist a mixed strategy equilibrium in which the expert is paid a fee

$$\alpha^D \triangleq \sqrt{\frac{\psi}{\hat{c}}\,(2\hat{c}\pi - \hat{\pi}\sigma_E^D)},$$

where

$$\sigma_E^D \triangleq \frac{\psi_H^a - \psi}{\psi_H^a - \psi_L^a} \in (0,1)$$

is the probability with which the expert gathers information, and the AA gathers information with probability

$$\sigma_A^D \triangleq \frac{1}{p}\left(1 - \frac{2\psi}{\alpha^D}\right).$$

This candidate equilibrium exists if $\alpha^D \in [\alpha_L, \alpha_H)$ (which can be verified by using numerical simulations) and it gives firms a higher expected profit than the one they would obtain by not hiring the expert, thus inducing the AA to bear the information acquisition cost. Clearly, as in the baseline model, the firms' expected profit from the mixed strategy equilibrium is decreasing in ψ and, when ψ is relatively large (i.e., it is close to ψ_H^a), the firms prefer not to hire the expert — i.e., the game under disclosure features the pure strategy equilibrium in which only the AA gathers information. The reason is simple: as

ψ grows larger, σ_E^D becomes smaller ($\sigma_E^D \to 0$ for $\psi \to \psi_H^a$) and obviously $\sigma_A^D < 1$, hence the probability with which the merger is approved is higher if the AA gathers information with certainty (and, clearly, in the latter case firms are better off also because they do not have to pay the expert).

It can be proved that, when p is too close to one, this equilibrium obtains for all $\psi \in (\psi_L^a, \psi_H^a)$. The reason is as follows: since in this model p denotes the *precision* of the AA's information acquisition technology only, the firms' incentives to free-ride on the AA are magnified when p grows larger. In conclusion, for an intermediate range of precision levels p, there exists a threshold $\overline{\psi}$ such that the mixed strategy equilibrium is played in the disclosure regime for $\psi \in (\psi_L^a, \overline{\psi})$.

In the regime without disclosure of the expert's contract, for every $\psi \in (\psi_L^a, \psi_H^a)$, there exists a mixed strategy equilibrium in which the expert is hired with probability

$$\sigma_F^S \triangleq \frac{\psi_H^a - \psi}{\psi_H^a - \psi_L^a} \in (0, 1)$$

and he is offered a bonus

$$\alpha^S \triangleq \frac{\sqrt{\psi^2 + 4\hat{c}(\hat{\pi} - (1-\hat{c})\psi)\psi)} - (1-2\hat{c})\psi}{2\hat{c}} \in (\alpha_L, \alpha_H),$$

whereas the AA gathers information with probability

$$\sigma_A^S \triangleq \frac{\hat{\pi} - \alpha^S}{p(\hat{\pi} - \alpha^S(1-2\hat{c}))} \in (0, 1).$$

The comparison between the two regimes from a consumer surplus perspective is easier when, for all $\psi \in (\psi_L^a, \psi_H^a)$, mixed strategies are played only in the game without disclosure of the expert's contract (i.e., when p is close to 1). In this case, clearly, the expected loss of consumer surplus is $-(1-p)\hat{c}w$ in the disclosure regime, whereas it is a decreasing function of ψ in the no disclosure regime. Moreover, using numerical examples, it can be shown that, when $\psi \to \psi_L^a$ (hence, $\sigma_F^S \to 1$), the regime without disclosure is the one which maximizes consumer surplus, unless p is too high. On the contrary, when $\psi \to \psi_H^a$, since, without disclosure, $\sigma_F^S \to 0$ and the AA gathers information with a probability lower than one, the disclosure of the expert's contract would be socially beneficial. As a consequence, there exists a unique threshold ψ^* such that the regime in which the expert's contract is disclosed to the AA increases consumer surplus relative to the regime without disclosure if and only if $\psi \geq \psi^*$. Our numerical simulations show that, as in the baseline model, this result holds true also when the mixed strategy equilibrium is played in the disclosure regime for $\psi \in (\psi_L^a, \overline{\psi})$. Interestingly, as illustrated in Figure 1.3 (which focuses on the region of parameters in which $\psi \in (\overline{\psi}, \psi_H^a))$, in this version of the model also the loss of consumer surplus in the no disclosure regime is, in absolute value, increasing in N, although less steep than the corresponding function in the disclosure regime. Nevertheless, in the next section, we provide numerical examples showing that ψ^* is still increasing in N, as in the baseline model. We can thus conclude that the results of Proposition 1.6 are robust when also type-II errors are taken into consideration.

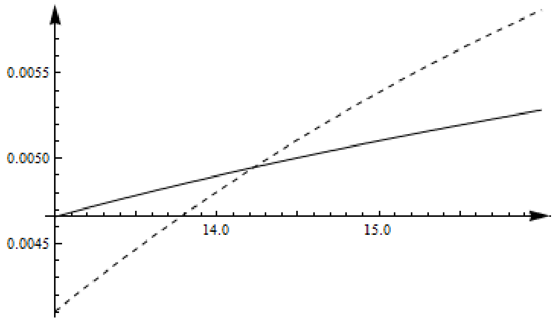

Fig. 1.3: Consumer expected loss (in absolute value) as function of N in the game with disclosure (dashed line) and without disclosure (continuous line), in a baseline example with $A = 10$, $p = 0.9$ and $\psi = 0.022$.

1.B.2. Additional computations

In what follows, we provide all the analytical details and some numerical results.

Preliminaries. Using Bayes rule, the posterior probability density functions are as follows

$$f(c|s=1) = \frac{1-c}{\int_0^1 (1-c)dc} = 2(1-c), \qquad f(c|s=0) = \frac{c}{\int_0^1 c\, dc} = 2c.$$

Accordingly,

$$\Pr[c < \hat{c}|s=0] = \int_0^{\hat{c}} f(c|s=0)dc = \hat{c}^2, \qquad \Pr[c > \hat{c}|s=1] = \int_{\hat{c}}^1 f(c|s=1)dc = 1-\hat{c}^2.$$

Define

$$\hat{w} \triangleq \int_0^1 w(c) f(c|s=1)dc =$$

$$= \frac{17 + 14N - 17N^2 - 8N^3 + 4A(-5 - 2N + 8N^2 + 2N^3) - 6A^2(2N^2-1)}{12N^2(1+N)^2}.$$

Then, the inequality $\hat{w} > 0$ is satisfied if and only if

$$N > \underline{N} \triangleq \frac{1}{24(A-1)} \left(12A^2 - 32A + 17 + \frac{144A^4 - 768A^3 + 1624A^2 - 1616A + 625}{\sqrt[3]{\kappa}} + \sqrt[3]{\kappa} \right)$$

where

$$k \triangleq 1728A^6 - 13824A^5 + 42480A^4 - 63008A^3 + 44484A^2 - 10680A - 1207+$$

$$+72\sqrt{2(A-1)^4\left(-1728A^6 + 16128A^5 - 62448A^4 + 128608A^3 - 148692A^2 + 91512A - 23407\right)}.$$

It can be easily checked that $\underline{N} \in (A-2, 2(A-2))$. Thus, in the remainder, we restrict our attention to $N \in [\underline{N}, 2(A-2)]$. Figure 1.4 shows that, in our baseline example, in which we set $A = 10$, we have to restrict our attention to $N \in [13, 16]$.

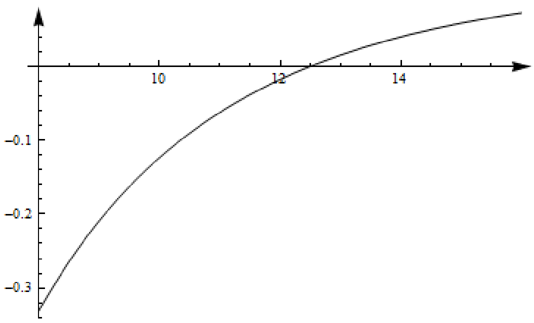

Fig. 1.4: $\hat{w}$ as function of N, in our baseline example, with $A = 10$.

Moreover, we have

$$\hat{\pi} = \int_0^1 \pi(c) f(c|s=1)dc = \frac{1}{2} + \frac{17 - 20A + 6A^2}{6N^2} + \frac{-7+4A}{3N} - \frac{2(A-1)^2}{(1+N)^2},$$

and it can be immediately verified that

$$\hat{\pi} > \hat{w} \iff 17+26N-23N^2-8N^3+6N^4+4A(-5-14N+2N^2+2N^3)+6A^2(1+4N) > 0,$$

which is clearly satisfied for all $N > 2$. Finally, we compute

$$w_I = \frac{(2-A+N)^3(2-A+N+4(A-1)N^2)}{12N^2(1+N)^4} > 0, \quad w_{II} = -\frac{(A-1)^4(4N^2-1)}{12N^2(1+N)^4} < 0.$$

We then turn to compare the four thresholds on the information acquisition cost parameter. It is trivial to see that, for all $p \in (0,1)$: $\psi_L^e < \psi_H^e$. To see that $\psi_L^a < \psi_H^a$, first it can be shown that

$$\psi_L^a = p\frac{(A-1)^6(4N^2-1) + (2-A+N)^5(2-A+N+4(A-1)N^2)}{24N^2(1+N)^6}$$

is decreasing in N, whereas we already proved that ψ_H^a is increasing in N. Then, it can be easily checked that $\psi_L^a|_{N=\underline{N}} < \psi_H^a|_{N=\underline{N}}$. Next, a sufficient condition for $\psi_H^a < \psi_H^e$ is

$w < \hat{\pi}$. We know that w is increasing in N. On the contrary,[34]

$$\frac{\partial \hat{\pi}}{\partial N} = \frac{1}{3}\left(\frac{12(A-1)^2}{(N+1)^3} - \frac{6A^2-20A+17}{N^3} - \frac{4A-7}{N^2}\right) < 0.$$

Thus, a sufficient condition for $\psi_H^a < \psi_H^e$ is $w|_{N=2(A-2)} < \hat{\pi}|_{N=2(A-2)}$, which can be easily verified. Finally,

$$\psi_L^e < \psi_L^a \iff \frac{(1-p)^2}{p(1-p(1-2\hat{c}))} < \frac{l}{\hat{\pi}}.$$

Notice that this inequality is (is not) satisfied at $p \to 1$ (resp., at $p \to 0$), its left-hand side is decreasing in p (and its right-hand side does not depend on p), and

$$\psi_L^e = \psi_L^a \iff p = \underline{p} \triangleq \frac{\hat{\pi}}{\hat{\pi} + l + \sqrt{l(4\hat{c}\hat{\pi} + l)}} \in (0,1).$$

Therefore, we can conclude that $\psi_L^e < \psi_L^a \iff p > \underline{p}$, which we assume throughout.

Disclosure regime. For $\psi \leq \psi_L^a$ and $\psi \geq \psi_H^a$, the SPNE of the game is as in the baseline model. For $\psi \in (\psi_L^a, \psi_H^a)$:

- if the expert is not hired, the AA gathers information, and the firms' expected profit is $p\hat{c}\pi$;
- if the expert is hired and offered a contingent fee α_H, he gathers information, whereas the AA does not, and the firms' expected profit is $\frac{1}{2}(\hat{\pi} - \alpha_H) = \frac{\hat{\pi}}{2} - \frac{\psi}{1-p}$;
- if the expert is hired and offered any contingent fee $\alpha \in [\alpha_L, \alpha_H)$, then, in the information acquisition game between the expert and the AA, there exists (besides the two pure strategy equilibria in which only one of the two players gather information) a mixed strategy equilibrium in which the expert gathers information with probability such as to make the AA indifferent between gathering information or not — i.e.,

$$\sigma_E^D \triangleq \frac{\psi_H^a - \psi}{\psi_H^a - \psi_L^a} \in (0,1)$$

[34] To establish this result, notice that

$$N > \underline{N} \iff A < \tilde{A} \triangleq \frac{-10 - 4N + 16N^2 + 4N^3 + \sqrt{2(N+1)^2\left(8N^4 + 2N^2 - 1\right)}}{12N^2 - 6},$$

and $\frac{\partial \hat{\pi}}{\partial N} < 0$ if and only if

$$A < \frac{2(N-2)(N+6)N^2 - 28N - 10 + \sqrt{2(N+1)^3(-1-2N+2N^2+26N^3-11N^4+2N^5)}}{6(-1-3N-3N^2+N^3)},$$

this threshold being higher than $\tilde{A}$.

and the AA gathers information with probability

$$\sigma_A^D(\alpha) \triangleq \frac{1}{p}\left(1 - \frac{2\psi}{\alpha}\right),$$

such as to make the expert (whose contingent fee is disclosed and given by α) indifferent between gathering information or not. If the firms expect this mixed strategy equilibrium to be played in the information acquisition game, then the (candidate) optimal contingent fee is determined by solving

$$\max_{\alpha\in[\alpha_L,\alpha_H)} \sigma_E^D\sigma_A^D\left(p\hat{c}(\pi-\alpha)+\frac{1-p}{2}(\hat{\pi}-\alpha)\right)+(1-\sigma_E^D)\sigma_A^D p\hat{c}(\pi-\alpha)+\sigma_E^D(1-\sigma_A^D)\frac{1}{2}(\hat{\pi}-\alpha),$$

whose solution gives α^D. In our numerical simulations, we verified that, under the restriction $p > \underline{p}$, $\alpha^D \in (\alpha_L, \alpha_H)$: see Figure 1.5 for an example.

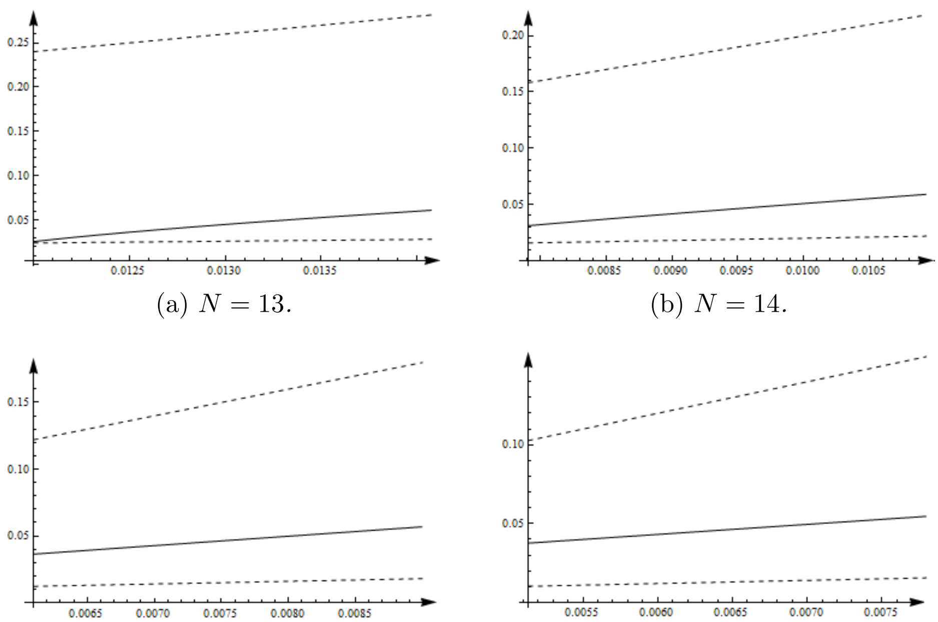

(a) $N = 13$. (b) $N = 14$.

(c) $N = 15$. (d) $N = 16$.

Fig. 1.5: Contingent fees α_L and α_H (dashed lines) and α^D (continuous line), in our baseline example with $A = 10$ and $p = 0.9$, as functions of $\psi \in (\psi_L^a, \overline{\psi})$.

We first compare the firms' payoffs in the two pure strategy equilibria. Since the firms' profit in the pure strategy equilibrium in which only the expert gathers information is decreasing in ψ, it follows that a sufficient condition in order for the other pure strategy equilibrium to be preferred is $p\hat{c}\pi > \frac{\hat{\pi}}{2} - \frac{\psi_L^a}{1-p}$, which is satisfied if p is sufficiently high — i.e.,

$$p > \tilde{p} \triangleq \frac{2\hat{c}\pi + \hat{\pi} + 2l - \sqrt{(2\hat{c}\pi + \hat{\pi} + 2l)^2 - 8\hat{c}\pi\hat{\pi}}}{4\hat{c}\pi}$$

Our numerical simulations show that this condition is always satisfied under the restriction $p > \underline{p}$ — i.e., $\tilde{p} < \underline{p}$: see Figure 1.6 for an example.
Therefore, in what follows, we compare the firms' profit when the mixed strategy equilibrium is played, which is given by

$$\hat{c}\left(\pi + 2\psi - 2\sqrt{2\pi\psi + \frac{\hat{\pi}\psi(\psi - \psi_H^a)}{c(\psi_H^a - \psi_L^a)}}\right) + \psi\frac{\psi - \psi_H^a}{\psi_H^a - \psi_L^a},$$

with their profit when the pure strategy equilibrium in which only the AA gathers information is played. For obvious reasons, the firms' profit in the mixed strategy equilibrium is decreasing ψ and, at $\psi = \psi_H^a$, firms are better off in the pure strategy equilibrium. On the contrary, at $\psi = \psi_L^a$, firms are better off in the mixed strategy equilibrium if and only if

$$p < \overline{p} \triangleq \frac{\hat{c}(\pi - 2\hat{\pi})l + \hat{c}^2\pi(\pi + 2l) - 2\sqrt{\hat{c}l(2\hat{c}\pi - \hat{\pi})(\hat{c}\pi^2 + (\pi - \hat{\pi})l)}}{(\hat{c}(\pi - 2l) + l)^2}.$$

In our numerical examples, we verified that $1 > \overline{p} > \underline{p}$ (see Figure 1.6 for an example). Thus, $p > \overline{p}$ is a sufficient condition in order for the pure strategy equilibrium in which only the AA gathers information to be the SPNE of the game with disclosure of the expert's contract for all $\psi \in (\psi_L^a, \psi_H^a)$.

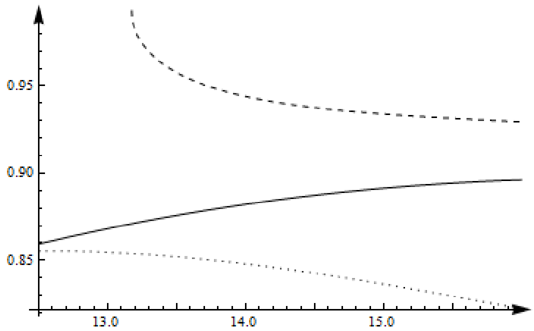

Fig. 1.6: Thresholds $\tilde{p}$ (dotted line), $\underline{p}$ (continuous line) and $\overline{p}$ (dashed line) as functions of N, in our baseline example, with $A = 10$.

For $p \in (\underline{p}, \overline{p})$, denote by $\overline{\psi}$ the threshold on the information acquisition cost such that the mixed strategy equilibrium is played for all $\psi \in (\psi_L^a, \overline{\psi})$, whereas the pure strategy equilibrium in which only the AA gathers information is played for all $\psi \in (\overline{\psi}, \psi_H^a)$. Table 1.1 shows the values for $\psi_L^a, \overline{\psi}$ and ψ_H^a in our baseline example.

No disclosure regime. The pure strategy equilibria are as in the baseline model. For every $\psi \in (\psi_L^a, \psi_H^a)$, there exists a mixed strategy equilibrium in which the firms hire the expert with a probability σ_F^S such as to make the AA indifferent between gathering information or not (implying $\sigma_F^S = \sigma_E^D \in (0, 1)$), and pay the expert a contingent fee α^S which makes

	ψ_L^a	$\overline{\psi}$	ψ^*	ψ_H^a
$N = 13$	.0111322	.0140740	.0202356	.0369391
$N = 14$	.0079119	.0109251	.0217215	.0432490
$N = 15$	.0060997	.0089942	.0228088	.0485055
$N = 16$	.0051336	.0077948	.0235926	.0528191

Table 1.1: Values for $\psi_L^a, \overline{\psi}, \psi^*$ and ψ_H^a in a baseline example with $A = 10$ and $p = 0.9$.

him indifferent between gathering information or not, given the AA's mixed strategy — i.e.,

$$\alpha^S = \frac{2\psi}{1 - p\sigma_A^S} \in (\alpha_L, \alpha_H).$$

The probability σ_A^S is in turn obtained by imposing that the firms are indifferent between hiring the expert or not — i.e.,

$$\sigma_A^S = \frac{\hat{\pi} - \alpha^S}{p(\hat{\pi} - \alpha^S(1 - 2\hat{c}))} \in (0, 1).$$

Solving these conditions yields the equilibrium values.

Optimal disclosure rule. If $p > \overline{p}$, then, for all $\psi \in (\psi_L^a, \psi_H^a)$, the relevant comparison is between the pure strategy equilibrium in which only the AA gathers information (played in the disclosure regime) and the mixed strategy equilibrium played if the expert's contract is not disclosed. In this case, the expected loss of consumer surplus does not depend on ψ in the disclosure regime, whereas it is a decreasing function of ψ in the no disclosure regime, and, for $\psi \to \psi_H^a$, this loss is (in absolute value) higher under no disclosure. Thus, there exists a threshold ψ^* such that the regime in which the expert's contract is disclosed to the AA increases consumer surplus relative to the regime without disclosure if and only if $\psi \geq \psi^*$ if and only if, when $\psi \to \psi_L^a$, the regime without disclosure is the one which maximizes consumer surplus. It can be easily seen that, for $\psi \to \psi_L^a$, the loss of consumer surplus in the no disclosure regime is

$$\mathcal{L}^S|_{\psi \to \psi_L^a} = -(1 - p\sigma_A^S|_{\psi \to \psi_L^a})l = -\frac{l}{\hat{\pi}}\left((1 - 2\hat{c})pl + \sqrt{pl(pl + 4\hat{c}(\hat{\pi} - (1 - \hat{c})pl))}\right),$$

which is clearly decreasing in p. On the contrary, $\mathcal{L}^D = -(1 - p)\hat{c}w$ is increasing in p. Moreover, it can be immediately seen that, when $\psi \to \psi_L^a$, the regime without disclosure is the one which maximizes consumer surplus if and only if p is not too large. Our numerical simulations show that this can be the case also for $p > \overline{p}$ (unless p gets very close to one), therefore it is always satisfied when $p \in (\underline{p}, \overline{p})$: see Figure 1.7 for an example.

As for this region of parameters, as in the baseline model, for $\psi \in (\psi_L^a, \overline{\psi})$ the relevant

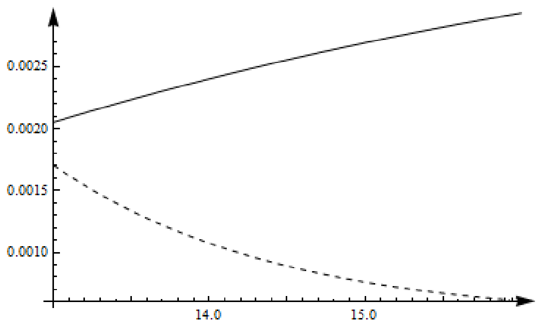

Fig. 1.7: Loss of consumer surplus (in absolute value), when $\psi \to \psi_L^a$, in the no disclosure regime (dashed line) and in the disclosure regime (continuous line), as functions of N, in our baseline example (with $A = 10$). We set $p = 0.95$, which is higher than $\overline{p}$ (see Figure 1.6).

comparison is between the two mixed strategy equilibria characterized above, whereas for $\psi \in (\overline{\psi}, \psi_H^a)$ we must compare the pure strategy equilibrium in which only the AA gathers information (arising in the disclosure regime) with the mixed strategy equilibrium played if the expert's contract is not disclosed.

For $\psi \in (\psi_L^a, \overline{\psi})$, since $\sigma_F^S = \sigma_E^D$, the regime without disclosure is preferable from the consumers' viewpoint if and only if $\sigma_A^S > \sigma_A^D$. Our numerical simulations show that this is the case for all $\psi \in (\psi_L^a, \overline{\psi})$: see Figure 1.8 for an example.

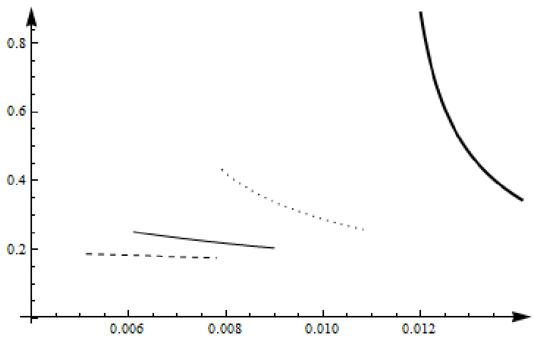

Fig. 1.8: Difference $\sigma_A^S - \sigma_A^D$ as function of $\psi \in (\psi_L^a, \overline{\psi})$, in a baseline example with $A = 10$ and $p = 0.9$, for $N = 13$ (thick line), $N = 14$ (dotted line), $N = 15$ (continuous line), $N = 16$ (dashed line).

As for $\psi \in (\psi_L^a, \overline{\psi})$, our numerical simulations show that $\psi^* > \overline{\psi}$ (see Table 1.1 for an example).

Finally, the results summarized in Table 1.1 suggest that the comparative statics of ψ^* with respect to N is as in the baseline model — i.e., ψ^* is increasing in N.

Chapter 2

When Prohibiting Platform Parity Agreements Harms Consumers

2.1 Introduction

Motivated by the recent antitrust scrutiny of price parity provisions, a number of contributions have examined the effects of such contractual agreements on firms' profits and consumer welfare (e.g., Edelman and Wright, 2015, Johansen and Vergé, 2016, among others). These models consider one or more competing sellers supplying products both through their own direct distribution channels and competing platforms, which are accessed by final consumers (see Figure 2.1 panel (a)).

Within this framework, two different types of contractual arrangements between a seller and a platform are considered, namely *narrow* and *wide price parity agreements*. Under a wide parity agreement, the price charged to final consumers in the direct distribution channel must not be lower than the price charged to final consumers through either of the indirect booking channels, not exceed the price charged in the direct distribution channel and, in addition, the prices charged through the two platforms must be identical. Instead, under a narrow parity agreement, the prices charged for products distributed through a certain platform may be different from the prices charged to consumers booking through another platform. However, as with the wide parity agreement, the price charged to final consumers in the direct distribution channel must not be lower than the price charged to final consumers through either of the indirect booking channels.

The literature concludes that typically wide parity agreements are anti-competitive, absent efficiencies, unless upstream competition is relatively fierce and sellers can delist from a platform.[1] The basic intuition behind this result is rather simple and is aligned with the theory of harm developed by several antitrust authorities (see, e.g., Akman, 2016).

[1]If a seller can exit one of the platforms and faces competition from other sellers in that platform, then it can leverage the possibility of delisting credibly and effectively, in which case wide parity may lead to lower commissions and final prices. Intuitively, this is because by exiting one platform a seller *de facto* reduces its marginal cost and hence can steal business from its competitors, whereas the sales lost are not particularly valuable if sellers are close competitors.

The considered agreements soften platform competition because a platform setting high commissions (or fees) will not lose market share since sellers cannot offer more favorable prices through alternative distribution channels, including the direct distribution channel, which may involve lower costs. Instead, the platform can charge high fees knowing that those fees will be spread across all transactions and that consumers will not be able to find lower-cost alternatives elsewhere.[2] Similar results apply to narrow parity agreements: when platforms are *must have*, platforms may not undercut each other even when possible. The reason is that any reduction in fees would not be compensated by an increase in sales, since sellers would have no incentive to reduce the price charged in the undercutting platform at the expense of their direct distribution channels (with a price tied to the price of the high commission platform) when direct and indirect distribution are close substitutes for a majority of final consumers.

These conclusions may well apply to industries characterized by the structure described above, such as the hotel booking industry, in which platforms compete with each other to attract final consumers. However, in several other cases (e.g., the airline ticket distribution industry) platforms are accessed by specialized intermediaries, which in turn are in competition with each other in the retail market (see Figure 2.1 panel (b)). In these industries, which involve vertical supply chains that are more complex than the hotel booking industry's, it arises a multiple marginalization problem which is different from the one studied in the available models: as in the literature, a seller will mark up the commissions charged by the platforms, which in turn will negotiate fees above their marginal costs. However, in addition to this, intermediaries will mark up the prices offered by the sellers, implying that final prices reflect two mark-ups. These differences in vertical structure are likely to have important economic implications as for the competitive effect of price parity clauses, also because the contractual provisions themselves are rather different with respect to the narrow and wide price parity agreements detailed above. Notably, in these industries the so called *platform parity provisions* are typically negotiated between sellers and platforms. Such clauses require sellers to provide the same products and related content (e.g., ancillary services) to intermediaries using one platform on no worse terms and conditions than the seller itself would apply to users of other platforms. The cumulative effect of such provisions in agreements of several platforms could lead to a situation where users of all platforms have access to the same content. However, these provisions do not

[2]In addition to this, wide parity agreements may also limit the entry and expansion of new platforms, and thus have a negative impact on investment and innovation (see Boik and Corts, 2016). Platforms that are not yet established will not be able to compete effectively in the supply of indirect distribution services by offering lower fees (in return for which they might negotiate lower prices with final consumers). Incumbent platforms will thus be able to capture an increasingly large share of consumer traffic, as network effects draw consumers and sellers to the most heavily used platforms.

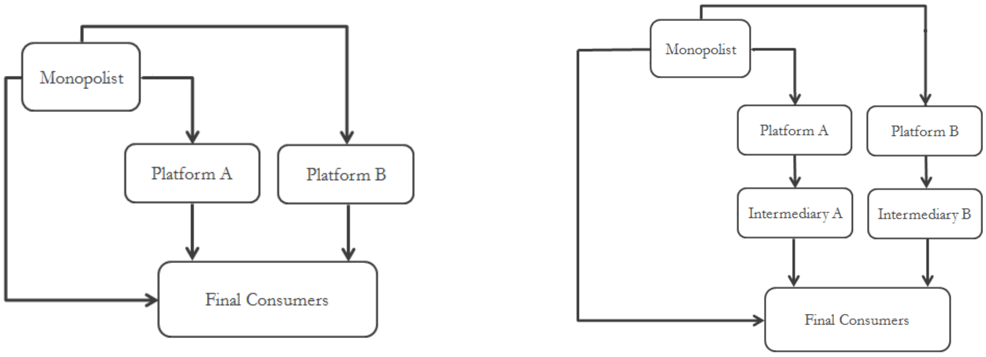

(a) *Vertical structures considered in the available models.*

(b) *Vertical structure considered in our model.*

Fig. 2.1: Industry vertical structure.

prevent sellers from distributing exclusive content through their *direct* distribution channels (e.g., their own websites). Furthermore, sellers are not required to ensure that final consumers pay the same prices in all *indirect* distribution channels (e.g., at intermediaries relying on different platforms). For these reasons, such agreements cannot be considered *price* parity agreements, as the wide and narrow parity agreements discussed above. The only operative constraint imposed by the parity agreements observed in the industries under consideration concerns the prices offered by a seller through the indirect distribution channels, which constitute the content that intermediaries distribute. Since these prices are forced to be identical, we may refer to this clause as a *platform (or content) parity agreement*.

We argue that the effects of such agreements on consumer welfare cannot be inferred from the analyses conducted in the available literature. The reasons are as follows. First, since, unlike the wide and narrow parity agreements reviewed above, prices set on the direct sale channels are unconstrained, each platform's incentive to negotiate high commissions or fees is limited, given that sellers will be able to increase the prices offered to the intermediaries that distribute their products indirectly and, therefore, divert business from the platform demanding high fees to their own direct distribution channels. Second, while platform parity may allow platforms to negotiate higher commissions, those extra rents are bound to be competed away because platforms compete with each other to increase their share of consumer traffic by expanding their network and supporting the competitive position of their intermediaries. Finally, since a seller must be concerned about the impact of multiple mark-ups on sales, platform parity may reduce its incentives to set high prices when distributing through platforms, thus mitigating the marginalization problem and leading to lower final prices and higher sales. This is because any increase in the prices offered through intermediaries using a given platform will lead to a parallel increase in the prices offered through the intermediaries dealing with the other platforms. These price increases will then be passed on to consumers in the form of higher retail prices and will cause a reduction in the demand served through the indirect distribution channel, which will be offset only in part by the increase in the demand served through the seller's direct distribution channel. Therefore, while platform parity may have similar rent shifting effects upstream to wide or narrow parity, its impact on consumer welfare is bound to be different.

To illustrate these points we build a three-level supply chain model where a monopolistic seller distributes its products both directly through its own distribution channel and indirectly through two platforms accessed by intermediaries, which in turn rely on the IT infrastructure provided by specialized platforms to buy the seller's product on behalf of final consumers (Figure 2.1 panel (b)). We assume that platforms and intermediaries are in exclusive relationships (e.g., because of switching costs) and, following Boik and Corts (2016) and industry practice, that contracts are linear — i.e., input prices do not vary with sales volumes.

As a benchmark, we first consider a wholesale industry where the seller sets the price on its direct channel and bilaterally negotiates a commission (wholesale price) with each platform for every unit of product purchased through that platform. Every platform, in turn, charges its own intermediary an 'access' price for each unit of product sold to final consumers. Finally, intermediaries set retail prices in competition between themselves and with the monopolist's direct distribution channel. In this setting, we find that a platform parity agreement is always pro-competitive even in the absence of efficiencies: the constraint on wholesale prices imposed through such a provision is used by the monopolist as a commitment device to mitigate the multiple marginalization problem. Indeed, when the platforms cannot be price discriminated, any attempt of the monopolist to increase the wholesale price charged to one platform immediately translates into a parallel shift of the wholesale price charged to the other platform. As a result, both platforms symmetrically increase the access price charged to the intermediaries, whereby leading to higher retail prices (since both intermediaries will mark up the increased access prices). By contrast, without platform parity, when the monopolist increases the wholesale price charged to a platform, the platform takes as given the wholesale price charged to its rival because contracts are secret and our solution concept is contract equilibrium. Hence, the multiple marginalization problem is relatively less important for the monopolist compared to the regime with platform parity. In other words, the excessive pass on rate that occurs under the parity provision refrains the monopolist from charging a wholesale price that is too high in equilibrium. This mandates a lower access price, a lower final price and thus a higher consumer surplus in the indirect distribution channel. In addition, since prices are strategic complements, lower prices in the indirect channel induce lower prices also in the direct channel, which benefits consumers in that segment too. Finally, by reducing the multiple marginalization problem, the provision also increases profits — i.e., the monopolist, the platforms and the intermediaries are better off with than without platform parity. Hence, in the wholesale model all players have aligned preferences.

Building on these insights we then consider an agency model (see, e.g., Johnson, 2017) which seems to reflect more closely industry practices. In this model each platform negotiates a per-unit commission with the monopolist for each unit of product purchased through that platform. The monopolist charges intermediaries an access price that they must pay for each unit of product purchased.[3] Intermediaries set retail prices (paid by final consumers) in competition between themselves and with the monopolist's direct distribution channel. In industries with such a vertical structure, we find that platform parity provisions might be pro-competitive (absent efficiencies) depending on the degree of product differentiation across distribution channels. Specifically, we show that the constraint on access prices implied by a platform parity agreement generates a new trade-off shaped by the following effects. First, each platform anticipates that, being concerned with double marginalization, under the parity provision the monopolist has a lower incentive to pass on commissions to the intermediaries. Hence, as in Boik and Corts (2016), platforms' fees are higher under the parity provision, which clearly harms consumers because it creates marginalization. Second, as in the wholesale model, the provision mitigates the marginalization problem between the monopolist and the intermediaries, which benefits consumers. Third, since the price in the direct channel is lower with platform parity than without platform parity (because the monopolist has an incentive to divert business towards that channel when the provision is in place), consumers on that segment benefit from the provision.

The net effect points in the direction of increasing consumer surplus when the products or the services provided through different distribution channels are not too differentiated. Essentially, in equilibrium, competition in the product market — i.e., within and between the distribution channels — erodes the intermediaries' mark ups and magnifies the pro-competitive effect of the provision on the marginalization problem. Notably, when platforms benefit from the provision consumers do too — i.e., their preferences are aligned with respect to the choice of the contractual provision. By contrast, absent efficiencies, the seller and the intermediaries are always better off without the provision. However, we also find that platforms can persuade more easily the intermediaries to prefer the parity regime through appropriate side payments than the monopolist. In other words, whenever

[3] We define the access price as the (unit) input price (net of commissions, surcharges and discounts) that an intermediary has to pay the monopolist in order to operate a transaction through a given platform.

the joint profit of the platforms and the intermediaries is higher with than without the parity provision, the total industry profit may well be lower without the provision — i.e., the price that the monopolist would require to accept the parity regime is higher than the gain obtained jointly by the platforms and the intermediaries. As a result, in these cases, the only way to increase consumer surplus is to allocate more decision rights to the platforms than the monopolist.

Summing up, based on our analysis we can conclude that, even with an upstream monopoly, content (i.e., platform) parity provisions cannot be presumed anti-competitive absent efficiencies. Interestingly, consumers and platforms' preferences are always aligned: as long as platforms benefit from platform parity, consumers gain as well (which is not always the case for the seller and the intermediaries). Hence, in practice, the likelihood that the introduction of such a provision benefits consumers is higher when platforms are not against it.

Finally, we develop some interesting extensions and robustness checks of the baseline model. First, we show that the welfare results discussed above hold qualitatively when considering multiple (more than two) competing platforms. Second, if the monopolist can commit to the price set on the direct channel before contracting with platforms, the pro-competitive effect of the provision survives only in the wholesale model. The reason why with commitment the provision is anticompetitive in the agency model is that the monopolist cannot compensate the effect of increased platforms' commissions with a lower price in the direct channel. This is because, being chosen at the outset of the game, that price is set 'efficiently' regardless of whether the provision is in place or not. Third, we show that the monopolist may use the quality of the product distributed through the indirect channel strategically in order to break down platform parity at the consumers' expense.

The remainder of the paper is organized as follows. After reviewing the related literature, in Section 2.2 we set-up the wholesale model. In Section 2.3, we analyze the agency model. Section 2.4 concludes. Proofs and some additional results are presented in the Appendix.

Related literature. Our paper contributes to the literature on wholesale and platform most-favored nation (MFN) clauses.[4] Early contributions (DeGraba and Postlewaite, 1992, McAfee and Schwartz, 1994, DeGraba, 1996, Marx and Shaffer, 2004) consider sequential contracting between a manufacturer and a number of retailers, and investigate the role played by MFN clauses in mitigating the time inconsistency problem faced by

[4]Other works examine MFN clauses that sellers offer to consumers: see, e.g., Cooper (1986), Butz (1990), Schnitzer (1994).

the supplier of a durable input. Following a number of antitrust investigations against the use of such clauses in wholesale markets (see, e.g., Avilés-Lucero and Boik, 2018) the potential pro- and anti-competitive effects of these provisions have been informally discussed in the law and economics literature (see, e.g., Baker and Chevalier, 2012).
More recent contributions investigate the welfare effects of the adoption of MFN clauses in online markets. Boik and Corts (2016), for example, consider a monopolist facing two competing platforms, which first simultaneously choose whether to impose a price parity agreement, then set per-unit commissions. After observing these choices, the monopolist sets final prices on both platforms (the so called agency model). Within this framework, a price parity clause is unambiguously anti-competitive since it raises platforms' commissions and retail prices. Similar results are found by Johnson (2017), who models competition in the upstream market. The anti-competitive nature of price parity provisions is challenged by Johansen and Vergé (2016), who consider endogenous platform participation and the presence of direct sales channels in addition to upstream competition (two ingredients that are also present in our model).[5] In their setting, if upstream competition is fierce enough, consumers benefit from the introduction of a narrow or a wide price parity clause provided that sellers can delist from platforms charging excessively high commissions. The mechanism through which parity agreements benefit consumers in our model is different from the argument put forward by Johansen and Vergé (2016), since we consider a monopolistic seller who never finds it optimal to delist from a platform.
Other contributions (Ronayne and Taylor, 2018, Calzada et al., 2018, Wang and White, 2016, Shen and White, 2019) suggest that platforms use these clauses to avoid *shoowrooming* — i.e., that consumers use the platform to learn of products, but then buy through the firms' direct sales channel if they find a lower price.[6] Edelman and Wright (2015), instead, assume that, as a result of costly investments, platforms are able to provide benefits to buyers,[7] and they show that a price parity clause leads to inflated retail prices, excessive adoption of the platforms' services, over-investment in benefits to buyers, and ultimately a reduction in consumer surplus. The welfare effects of price parity clauses in

[5]Actually, because of preference for variety, in our model delisting never occurs in equilibrium.

[6]In a model with a monopolistic platform, Ronayne and Taylor (2018) show that price parity agreements are profitable for the platform and reduce consumer surplus, whereas Calzada et al. (2018) find that these clauses may induce single homing by sellers, thereby reducing the products offered on each platform. Similar results are found by Wang and White (2016), who assume that the platform reduces consumers' search costs. However, when considering competing vertically differentiated platforms, they show that, under some circumstances, a narrow price parity agreement can enhance consumer surplus. Finally, Shen and White (2019) assume that the platform can make a recommendation to each consumer about which product to buy, finding that the main effect of price parity agreements is to shift surplus from sellers to the platform, and that these clauses can increase total welfare.

[7]These benefits can consist in offering complementary products, reducing transaction costs, and offering financial rebates. Moreover, consumers incur in a cost to join a platform and (like in Wang and White, 2016) the platform charges commissions to sellers as well as buyers.

the agency model are also discussed in the law and economics literature (e.g., Ezrachi, 2015). In all these models, the pro-competitive effect of price parity agreements is driven by the presence of efficiencies, which are instead absent in our model.
Finally, as for the comparison between the wholesale and the agency model, Foros et al. (2017) show that, even if platforms' commission rates remain the same across the two business models (for exogenous reasons), the use of price parity clauses may facilitate the adoption of the agency model which, in turn, may involve lower consumer prices. Our analysis shows that, like also in Johnson (2017), the agency model increases consumer surplus and platforms' profits when contracts and parity provisions are endogenously chosen within each business model.

2.2 The wholesale benchmark

In order to highlight the beneficial effects of platform parity agreements in the clearest possible way, it is useful to start with the standard wholesale framework, where these provisions unambiguously increase consumer welfare and firms' profits. We will then turn to the agency model and show that, in industries with such a business structure, platform parity agreements increase welfare as long as competition within and between the distribution channels is fierce enough — i.e., when the products or the services sold in these market segments are not too differentiated.

Markets and players. Consider a multi-channel and multi-tier industry in which a monopolist (M) sells its product through a direct channel and two competing platforms (each denoted by P_i, with $i = A, B$) accessed by intermediaries competing to attract final consumers. Suppose, for simplicity, that intermediaries and platforms are in exclusive relationships — e.g., because of switching costs.[8] Final consumers can buy the monopolist's product either through the direct sale channel or through the intermediaries.

The monopolist. The monopolist sets a price p_d for sales through its direct distribution system. Moreover, it can also sign distribution contracts with the platforms. Following the literature (see, e.g., Boik and Corts, 2016, and Gaudin, 2019), we assume linear contracts: each contract specifies a unit (wholesale) price t_i for every purchase processed through platform i. Production costs are assumed to be linear and normalized to zero without

[8]We assume exclusivity in order to understand the role of platforms in multi-channel and multi-layer markets. In fact, in the opposite scenario where platforms compete to attract intermediaries and are homogeneous, they would make zero profits in equilibrium and play no role in the analysis.

loss of generality. The monopolist's profit is

$$\pi^M(\cdot) \triangleq \sum_{i=A,B} t_i q_i + p_d q_d,$$

with q_i and q_d denoting the quantities sold through platform $i = A, B$ and the direct distribution system, respectively.

Platforms. Platforms also use linear contracts when dealing with intermediaries. Hence, each platform $i = A, B$ charges a unit (access) price w_i to its exclusive intermediary. Accordingly (normalizing to zero their costs) each platform's profit is

$$\pi_i^P(\cdot) \triangleq (w_i - t_i) q_i.$$

The platforms' outside option is normalized to zero without loss of generality. Hence, in this business model platforms act as wholesale distributors, aggregating many dispersed retailers. Their (implicit) role is to help sellers to reach out retailers and save on transaction costs.

Intermediaries. Each intermediary I_i charges a retail price p_i at which final consumers can buy M's product. Hence, I_i's profit is

$$\pi_i^I(\cdot) \triangleq (p_i - w_i) q_i,$$

where, for simplicity, the distribution cost is equal to zero. Again, we normalize to zero (without loss of generality) the intermediaries' outside option.

Demand functions. In line with the empirical evidence (see, e.g., Cazaubiel et al., 2018) we assume that consumers perceive the products sold through the direct and the indirect channel as imperfect substitutes. Some consumers may in fact prefer to purchase through the indirect channel because intermediaries offer (un-modelled) additional services that M is unable or unwilling to supply in the direct channel.
To this purpose, following the literature (e.g., Johansen and Vergé, 2016), for the sake of tractability, we assume that the demand system reflects the preferences of a representative consumer, whose utility function is

$$U(\cdot) \triangleq \sum_{j=A,B,d} q_j - \frac{1}{2} \sum_{j=A,B,d} q_j^2 - \gamma \sum_{i,j=A,B,d;,j\neq i} q_j q_i - \sum_{j=A,B,d} p_j q_j + m, \tag{2.1}$$

where m is the utility from income. Notice that since this utility function is strictly concave, it displays preference for variety. This means that the representative consumer prefers to have more consumption options.[9] As we will explain, this property will have important consequences on our analysis.

Standard techniques then yield the (direct) demand functions

$$q_A \triangleq D^A(p_A, p_B, p_d) = \frac{1-\gamma-(1+\gamma)\,p_A+\gamma\,(p_B+p_d)}{(1-\gamma)(1+2\gamma)},$$
$$q_B \triangleq D^B(p_B, p_A, p_d) = \frac{1-\gamma-(1+\gamma)\,p_B+\gamma\,(p_A+p_d)}{(1-\gamma)(1+2\gamma)},$$

and

$$q_d \triangleq D^d(p_d, p_A, p_B) = \frac{1-\gamma-(1+\gamma)\,p_d+\gamma\,(p_A+p_B)}{(1-\gamma)(1+2\gamma)}.$$

Hence, γ reflects the degree of substitutability between products within and across distribution channels. In the Appendix we show that the analysis' results do not change qualitatively if we assume that the products distributed in the indirect channel are perceived by the consumer as closer substitutes compared to the product distributed in the direct channel. We assume that $\gamma \in [0, \overline{\gamma}]$, with $\overline{\gamma} \approx 0.9$ in order to guarantee that second-order conditions hold (see the Appendix).

Timing. The timing of the game is as follows (see Figure 2.2):

$t = 1$ M simultaneously offers contracts — i.e., t_A and t_B.

$t = 2$ Upon observing its own offer t_i, platform P_i sets w_i.

$t = 3$ Upon observing its own offer w_i, intermediary I_i sets p_i and M sets p_d, simultaneously.

$t = 4$ After observing the vector of prices $\mathbf{p} \triangleq (p_d, p_A, p_B)$, final consumers allocate their demand.

Platform (content) parity. We consider two versions of the game, depending on the presence or not of a *platform (content) parity* agreement between the monopolist and the platforms. Specifically:

- When a *platform parity* agreement is in place, the monopolist commits to offer the same wholesale price to the competing platforms — i.e., $t_A = t_B = t$.

[9] For example, in the airline ticket distribution industry, the consumer may want to have the possibility to book in different ways because it may not always be able to reach directly the airline website.

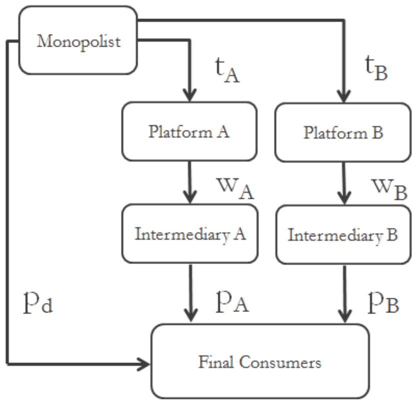

Fig. 2.2: The wholesale model.

- Without such agreement, the monopolist is free to offer different prices $t_A \neq t_B$ to the platforms.

Equilibrium concept. For the sake of tractability, as in Johansen and Vergé (2016) and Rey and Vergé (2017), our solution concept will be *Contract Equilibrium* (see Crémer and Riordan, 1987, and Horn and Wolinsky, 1988). We focus on symmetric equilibria in which, depending on whether a platform parity agreement is in place ($k = 1$) or not ($k = 0$):[10]

- M sets $p^*_{d,k}$ and charges the same wholesale price t^*_k to both platforms.
- Both platforms charge the same access price w^*_k to the intermediaries.
- Each intermediary charges the same retail price p^*_k to final consumers.

As noted by Rey and Vergé (2017), this equilibrium concept has some of the features of a perfect Bayesian Nash equilibrium with passive beliefs: P_i chooses the access price w_i charged to I_i assuming that its rival remains under the equilibrium contract t^*_k, even if P_i has received an out-of-equilibrium contract $t_i \neq t^*_k$ from M. This is in line with the market-by-market bargaining restriction of Hart and Tirole (1990) and with the passive beliefs or pairwise-proofness assumption of McAfee and Schwartz (1994). Similarly, at the stage in which retail prices are set, I_i chooses p_i assuming that its rival remains under the equilibrium contract w^*_k, even if I_i has received an out-of-equilibrium contract $w_i \neq w^*_k$ from P_i. Hence, the equilibrium concept that we use discards the possibility of multilateral

[10]It can be easily proved that, in our linear demand setting, there is a unique Contract Equilibrium and that this equilibrium is symmetric. However, to simplify exposition, we directly focus on a symmetric equilibrium.

deviations by the monopolist, who is the only player making multiple offers in the game. We discuss alternative equilibrium concepts in Section 2.3.8.

Multi-product monopolist. To begin with, it is useful to characterize the solution of the benchmark in which M sells directly to all consumers — i.e., the outcome of the game in which M is vertically integrated with the platforms and the intermediaries. In this hypothetical scenario, M's maximization problem is

$$\max_{p_A, p_B, p_d} \sum_{i=A,B} p_i D^i(p_i, p_{-i}, p_d) + p_d D^d(p_d, p_i, p_{-i}).$$

It can be shown (see the Appendix) that M charges the same price for all products

$$p^M = \frac{1}{2},$$

and that efficient quantities are symmetric since the utility function (2.1) displays preferences for variety — i.e.,

$$q_A^M = q_B^M = q_d^M = \frac{1}{2+4\gamma}.$$

In short, when M behaves as a multi-product monopolist, it fully internalizes the effects of intra- and inter-channel competition. In the rest of the analysis we will often refer to this benchmark in order to identify the extent of multiple marginalization.

2.2.1 Equilibrium analysis

In what follows we characterize the equilibrium of the game with and without platform parity. We will then study the impact of the parity provision on consumer surplus and firms' profit.

First we characterize the intermediaries' pricing behavior. Consider a symmetric equilibrium in which intermediaries charge p_k^*, platforms charge w_k^* and the monopolist charges t_k^* and $p_{d,k}^*$ on the direct channel. Hence, for any $k = 0, 1$ and for every offer w_i (received from P_i), I_i solves the following maximization problem

$$\max_{p_i} D^i\left(p_i, p_k^*, p_{d,k}^*\right)(p_i - w_i).$$

whose standard first-order condition is

$$\frac{1-\gamma-(1+\gamma)\,p_i+\gamma\left(p_k^*+p_{d,k}^*\right)}{(1-\gamma)(1+2\gamma)} - (p_i - w_i)\,\frac{1+\gamma}{(1-\gamma)(1+2\gamma)} = 0. \tag{2.2}$$

Given the equilibrium candidate under consideration, this condition defines I_i's 'best reply' to w_i — i.e.,

$$p_k\left(w_i\right) \triangleq \frac{w_i}{2} + \frac{1-\gamma+\gamma\left(p_k^* + p_{d,k}^*\right)}{2\left(1+\gamma\right)}, \quad \forall k = 1, 0. \tag{2.3}$$

As intuition suggests, this expression is increasing in the access price w_i and (since prices are strategic complements) in p_k^* and $p_{d,k}^*$ — i.e., the rivals' equilibrium prices in the indirect and direct channels, respectively.

Equilibrium with platform parity

When a platform parity agreement is in place, M charges the same wholesale price (t) to both platforms. Hence, P_i solves the following maximization problem

$$\max_{w_i} D^i\left(p_1\left(w_i\right), p_1^*, p_{d,1}^*\right)\left(w_i - t\right),$$

whose (standard) first-order condition yields

$$\frac{1-\gamma-(1+\gamma)\,p_1\left(w_i\right)+\gamma\left(p_1^*+p_{d,1}^*\right)}{(1-\gamma)(1+2\gamma)} - \frac{1+\gamma}{(1-\gamma)(1+2\gamma)}\frac{\partial p_1\left(w_i\right)}{\partial w_i}\left(w_i - t\right) = 0.$$

Using (2.2) yields

$$-\left(p_1\left(w_i\right) - w_i\right) + \frac{\partial p_1\left(w_i\right)}{\partial w_i}\left(w_i - t\right) = 0,$$

whose solution is

$$w_1^*(t) \triangleq \frac{t}{2} + \frac{1-\gamma+\gamma\left(p_1^* + p_{d,1}^*\right)}{2\left(1+\gamma\right)}. \tag{2.4}$$

This function is symmetric — i.e., it is the same for both platforms — and, as intuition suggests, it is increasing in the (retail) prices charged for the rival products (p_1^* and $p_{d,1}^*$). Next, substituting (2.4) into (2.3) and solving for p_1^*, we obtain:

$$p_1^*(t) \triangleq \frac{1+\gamma}{4+\gamma}t + \frac{3(1-\gamma+\gamma p_{d,1}^*)}{4+\gamma}.$$

Notice that $p_1^*(t)$ is increasing in t. When M charges a higher wholesale price to the platforms, they increase the access prices charged to the intermediaries, which increases the retail price charged for their product: a multiple marginalization effect.

Finally, we can examine M's maximization problem(s). At the final (pricing) stage M solves

$$\max_{p_d} 2tD^i\left(p_1^*\left(t\right), p_1^*\left(t\right), p_d\right) + p_d D^d\left(p_d, p_1^*\left(t\right), p_1^*\left(t\right)\right),$$

whose first-order condition is

$$\underbrace{\frac{1-\gamma-(1+\gamma)\,p_d+2\gamma p_1^*\,(t)}{(1-\gamma)(1+2\gamma)}-p_d\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}}_{\text{Monopoly rule}}+\underbrace{\frac{2\gamma t}{(1-\gamma)(1+2\gamma)}}_{\text{Horizontal externality}}=0. \tag{2.5}$$

The first term in this expression reflects the standard monopoly trade-off: increasing p_d lowers demand in the direct market but it increases M's profit on each unit of sale. The second term, instead, captures the externality that a higher price in the direct channel creates on the indirect one: a higher p_d increases the intermediaries' demand and thus (for given wholesale prices) the revenue that M collects from the platforms.
Letting $p_{d,1}^*\,(t)$ denote the solution of (2.5), M's maximization problem at stage 1 is

$$\max_t 2tD^i\left(p_1^*\,(t)\,,p_1^*\,(t)\,,p_{d,1}^*\,(t)\right)+p_{d,1}^*\,(t)\,D^d\left(p_{d,1}^*\,(t)\,,p_1^*\,(t)\,,p_1^*\,(t)\right).$$

Differentiating with respect to t, by the Envelope Theorem we obtain the following first-order condition

$$\underbrace{\frac{2\gamma t}{(1-\gamma)(1+2\gamma)}\frac{\partial p_1^*\,(t)}{\partial t}+\frac{1-\gamma-p_1^*\,(t)+\gamma p_{d,1}^*\,(t)}{(1-\gamma)(1+2\gamma)}}_{\text{Vertical externality}}+\underbrace{\frac{\gamma p_{d,1}^*\,(t)}{(1-\gamma)(1+2\gamma)}\frac{\partial p_1^*\,(t)}{\partial t}}_{\text{Horizontal externality}}=0.$$

Once again, this condition features two terms reflecting the impact of a higher t on M's total profit. First, other things being equal, a higher t increases the revenues from the indirect channel; but, since this also increases the (equilibrium) retail price in that channel, demand drops and so does M's revenue. Second, by increasing $p_1^*\,(t)$, a higher t also increases demand on the direct channel.
Imposing symmetry and solving the first-order conditions derived above, we can state the following.

Proposition 2.1. *With platform parity, the symmetric equilibrium of the wholesale model has the following features:*

(*i*) *The monopolist sets*

$$p_{d,1}^*=p^M+\underbrace{\frac{\gamma}{(1+2\gamma)\,(4+\gamma)}}_{\textit{Channel externality}},$$

and charges

$$t_1^*=p^M+\frac{3\gamma^2}{2\,(1+2\gamma)\,(4+\gamma)}.$$

(*ii*) *The platforms charge*

$$w_1^* = t_1^* + \underbrace{\frac{1-\gamma}{4+\gamma}}_{P_i\text{'s mark-up}}.$$

(*iii*) *The intermediaries set*

$$p_1^* = w_1^* + \underbrace{\frac{1-\gamma}{2\left(4+\gamma\right)}}_{I_i\text{'s mark-up}},$$

with $p_1^* \geq p_{d,1}^*$.

In the equilibrium every player makes a positive profit — i.e., $p_1^* \geq w_1^* \geq t_1^* > 0$. The reason is rather intuitive: under linear contracts, both the platforms and the intermediaries pass on their 'input prices' (the wholesale and the access price respectively) in order to secure positive margins. As a result, retail prices are higher than the price that would be charged by a multi-product monopolist: a multiple marginalization effect.

Equilibrium without platform parity

Next, we analyze the regime without platform parity, where M can charge different wholesale prices to the platforms. P_i solves the following maximization problem

$$\max_{w_i} D^i \left(p_0\left(w_i\right), p_0^*, p_{d,0}^*\right)\left(w_i - t_i\right),$$

whose first-order condition, using (2.2), is

$$-\left(p_0\left(w_i\right) - w_i\right) + \frac{\partial p_0\left(w_i\right)}{\partial w_i}\left(w_i - t_i\right) = 0,$$

yielding P_i's best reply to every t_i chosen by M — i.e.,

$$w_0^*(t_i) \triangleq \frac{t_i}{2} + \frac{1-\gamma+\gamma\left(p_0^* + p_{d,0}^*\right)}{2\left(1+\gamma\right)}.$$

Substituting $w_0(t_i)$ into $p_0(w_i)$ we obtain I_i's retail price as a function of t_i — i.e.,

$$p_0^*(t_i) \triangleq \frac{t_i}{4} + \frac{3\left(1-\gamma+\gamma(p_0^* + p_{d,0}^*)\right)}{4\left(1+\gamma\right)},$$

which, as expected, is increasing in t_i (the marginal cost) and the equilibrium prices charged in the final market — i.e., p_0^* and $p_{d,0}^*$.

Consider M's behavior. For every pair of contracts $\mathbf{t} \triangleq (t_i, t_{-i})$ offered to the platforms, M chooses p_d in order to solve

$$\max_{p_d} \sum\nolimits_{i=A,B} t_i D^i(p_0^*(t_i), p_0^*(t_{-i}), p_d) + p_d D^d(p_d, p_0^*(t_A), p_0^*(t_B)),$$

whose first-order condition is

$$\underbrace{\frac{1-\gamma-(1+\gamma)p_d+\gamma\sum_{i=A,B} p_0^*(t_i)}{(1-\gamma)(1+2\gamma)} - p_d\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}}_{\text{Monopoly rule}} + \underbrace{\frac{\gamma\sum_{i=A,B} t_i}{(1-\gamma)(1+2\gamma)}}_{\text{Horizontal externality}} = 0. \quad (2.6)$$

As before, at the (final) pricing stage M must take into account the positive effect of increasing p_d on the intermediaries' demand, and thus on its revenue from the indirect channel.

Let $p_{d,0}^*(\mathbf{t})$ be the solution of (2.6). Moving backward at the contracting stage, M chooses the wholesale price charged to P_i in order to solve

$$\max_{t_i} p_{d,0}^*(\mathbf{t}) D^d(p_{d,0}^*(\mathbf{t}), p_0^*(t_i), p^*) + t_0^* D^{-i}(p_0^*, p_0^*(t_i), p_{d,0}^*(\mathbf{t})) + t_i D^i(p_0^*(t_i), p_0^*, p_{d,0}^*(\mathbf{t})),$$

whose first-order condition (by using the Envelope Theorem) is

$$\underbrace{\frac{1-\gamma-(1+\gamma)p_0^*(t_i)+\gamma\left(p_0^*+p_{d,0}^*(\mathbf{t})\right)}{(1-\gamma)(1+2\gamma)} - t_i\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}\frac{\partial p_0^*(t_i)}{\partial t_i}}_{\text{Monopoly rule}} +$$

$$+\underbrace{\frac{\gamma}{(1-\gamma)(1+2\gamma)}\frac{\partial p_0^*(t_i)}{\partial t_i}\left(p_{d,0}^*(\mathbf{t})+t_0^*\right)}_{\text{Channel externality}} = 0.$$

Since we are considering a contract equilibrium, in the absence of a platform parity agreement, the monopolist itself, when contracting with P_i, takes as given the equilibrium contract offered to the other platform (and the resulting equilibrium price p_0^*). Hence, increasing t_i has three main effects on M's profit. First, it increases M's demand on the direct channel because $p_0^*(t_i)$ is increasing. Second, a higher t_i also increases the revenue that M obtains from P_{-i} because (*ceteris paribus*) it increases I_{-i}'s demand. Both effects create a channel externality which adds to the third more standard effect: a higher t_i increases the revenue that M collects from P_i, but it also lowers the demand for the product distributed through P_i since $p_0^*(t_i)$ is increasing.

Imposing symmetry and solving the first-order conditions derived above, we can state the following.

Proposition 2.2. *Without platform parity, the symmetric equilibrium of the wholesale model has the following features:*

(i) The monopolist sets

$$p^*_{d,0} = p^M + \underbrace{\frac{3\gamma\,(1+\gamma)}{(1+2\gamma)\,(8+5\gamma)}}_{\textit{Channel externality}},$$

and charges

$$t^*_0 = p^M + \frac{3\gamma\,(1+3\gamma)}{2\,(1+2\gamma)\,(8+5\gamma)}.$$

(ii) The platforms charge

$$w^*_0 = t^*_0 + \underbrace{\frac{2\,(1-\gamma)}{8+5\gamma}}_{P_i\textit{'s mark-up}}.$$

(iii) The intermediaries set

$$p^*_0 = w^*_0 + \underbrace{\frac{1-\gamma}{8+5\gamma}}_{I_i\textit{'s mark-up}},$$

with $p^*_0 > p^*_{d,0}$.

The equilibrium features multiple marginalization also in the regime without parity — i.e., $p^*_0 \geq w^*_0 \geq t^*_0 > 0$. The reason is as before: each level of the supply chain creates a mark up, the sum is ultimately passed on to final consumers.

2.2.2 Welfare

We can now examine the welfare effects of platform parity. Before stating the main result of the section, it is useful to observe that, when the parity rule is in place, the final price in the indirect channel is more sensitive to the wholesale price charged by the monopolist to the platforms — i.e.,

$$\frac{\partial p^*_1(t)}{\partial t} - \frac{\partial p^*_0(t_i)}{\partial t_i} = \frac{3\gamma}{4\,(\gamma+4)} > 0.$$

In other words, the rate at which intermediaries pass on wholesale prices to final consumers is higher when the provision is in place. In this regime, a platform (say P_i) that is charged a higher wholesale price anticipates that the other platform (say P_{-i}) faces the same price increase. Hence, when P_i observes a higher t it expects P_{-i} to pass on this higher wholesale price to its intermediary via a higher access price w_{-i}, which will in turn induce I_{-i} to increase the final price p_{-i}. As a result, when t increases, P_i has two reasons for charging I_i a higher access price w_i: first, because it faces a higher wholesale price and needs to mark up more its own intermediary: a standard vertical externality; second, because it expects a higher final price by I_{-i} and, as a result, a higher demand for I_i's product: a horizontal externality introduced by platform parity. By contrast, when the provision is not in place, P_i only observes its own wholesale price and, in the contract equilibrium, it assumes that the rival remains under the equilibrium contract.

Interestingly, the gap between the rate at which intermediaries pass on with and without parity is increasing in γ. Essentially, with parity an increase in t is equivalent to a common cost shock, which is passed on to a greater extent the more competition there is (the greater γ is). By contrast, without parity, a higher t_i is equivalent to an idiosyncratic cost shock, and (with linear demand) the pass on rate does not depend on γ.

Hence, when products become closer substitutes, the intermediaries are more responsive to the prices of their rivals because competition is more intense. Thus, the horizontal externality introduced by the parity provision is relatively more pronounced when γ is large. By contrast, when γ is small, intermediaries care less about the price of their rivals. In this case, the dominating force is the vertical externality — e.g., in the limit case with $\gamma \to 0$ the pass on rate is nearly the same regardless of whether there is parity or not. We can thus state the following.

Proposition 2.3. *In the wholesale model the introduction of a platform parity agreement lowers prices at every level of the supply chain — i.e., $t_1^* < t_0^*$, $w_1^* < w_0^*$, $p_{d,1}^* < p_{d,0}^*$ and $p_1^* < p_0^*$ — and thus it always benefits consumers. Moreover, it also increases firms' individual profits. Hence, in the wholesale model platform parity is total welfare enhancing.*

Hence, the provision can be interpreted as a commitment device that forces the monopolist to mitigate the multiple marginalization effect by choosing a lower wholesale price. The intuition hinges on the effect discussed above. Platform parity is equivalent to introducing a common component into the platforms and the intermediaries' cost function: an increase of such component has a symmetric impact on the pass on rates charged at both levels of the supply chain. In other words, when platforms cannot be price discriminated, any attempt of the monopolist to increase the (wholesale) price charged to one platform immediately translates into a parallel shift of the (wholesale) price charged to the other platform. As a result, both platforms increase the access price charged to their intermediaries, whereby leading to higher retail prices. By contrast, without platform parity, it is as if the cost function of the competing platforms and intermediaries are affected by idiosyncratic components only. In fact, contract equilibrium implies that when the monopolist increases the (wholesale) price charged to a platform (say P_i), this platform takes as given the access price charged by its rival since (as for the logic of passive beliefs) it remains under the equilibrium contract t_0^*. Therefore, a higher t_i only translates into a higher retail price charged by I_i, which makes the multiple marginalization problem less problematic for M. Precisely the increased pass on rate that occurs under the parity provision refrains M from charging a wholesale price that is too high in equilibrium. This mandates a lower access price, a lower final price and thus a higher consumer surplus in the indirect distribution channel. In addition, since prices are strategic complements, this effect propagates to the direct channel, leading to lower final prices in that segment also. As a result, also consumers in the direct channel benefit from the parity provision.

Figure 2.3 shows the difference between consumer surplus with and without the parity provision. It can be seen that, as γ increases, the positive impact of the provision on consumer surplus grows larger: as competition in the final market becomes more intense, the difference between the intermediaries' pass on rate with and without platform parity increases, which strengthens the beneficial commitment effect of the provision.

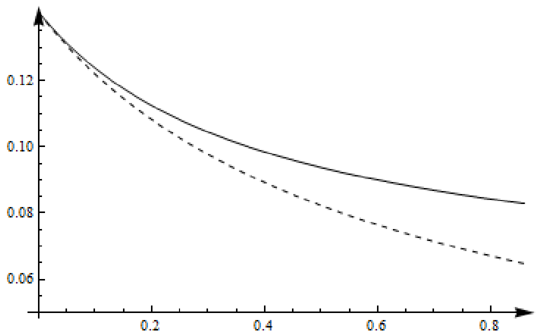

Fig. 2.3: Consumer surplus with platform parity (continuous line) and without platform parity (dashed line) as functions of γ.

Notably, platform parity echoes the effect of assuming symmetric beliefs off the equilibrium path when using PBE as solution concept (see, e.g., Rey and Tirole, 2007, and Pagnozzi and Piccolo, 2011). Essentially, as in our analysis, with symmetric beliefs a platform receiving from the monopolist an offer different from what it expects in equilibrium, believes that the competing platform has received the same offer, which generates welfare effects equivalent to those discussed above. Yet, in contrast to the previous literature, in our three layer model where the monopolist can also sell directly to consumers and under linear contracts, the effect on consumer surplus is pro-competitive rather than anti-competitive.

Finally, it should also be noted that, by reducing multiple marginalization, the provision increases profits at every level of the supply chain. This result is in line with Hart and

Tirole (1990) and McAfee and Schwartz (1994) who consider symmetric beliefs in a model with a single (monopolistic) manufacturer and two independent and competing retailers. They show that, with private contracts, the manufacturer's profit is higher with symmetric than with passive beliefs.

2.2.3 Discussion

Before turning to the agency model, a few remarks are in order.

Two-part tariffs. The effects highlighted above hinge on the multiple marginalization problem that is created by the assumption of linear contracts. One may wonder what would happen with two-part tariffs. In the wholesale model the answer is simple. When the monopolist can charge a fixed fee to the platforms, who can in turn charge a fixed fee to the intermediaries, the double marginalization problem wipes out. The reason is that the monopolist internalizes the profits of the entire indirect channel via the fixed fee charged to the platforms. This is because, by doing so the monopolist is able to internalize the fixed fee charged by the platforms to the intermediaries, and thus the intermediaries' profit. However, the monopolist cannot reach the efficient solution p^M because intermediaries compete downstream: conditional on wholesale prices being equal to marginal costs, prices are too low compared to what M would like to choose. In this case, platform parity can help the monopolist in restoring its monopoly power by increasing symmetrically the platforms' wholesale prices so to relax competition in the indirect channel. As a result, with two-part tariffs, platform parity is bound to be anti-competitive. Yet, when extracting profits by means of a fixed fee is costly (e.g., because of frictions akin moral hazard or adverse selection) our conclusions still hold as long as the cost of extracting profits up-front is sufficiently high (a similar point is discussed in Rey and Vergé, 2017).[11]

RPM. Up until now we have assumed that M cannot control the retail prices charged by the intermediaries. When M can dictate the retail prices in the indirect channel — i.e., when Resale Price Maintenance (henceforth, RPM) is allowed — the multiple marginalization problem wipes out (see, Motta, 2004, Ch. 6, for a survey). Hence, consumer surplus and M's profit increase, whereas the intermediaries and the platforms make zero profit. Obviously, in this case, platform parity is always welfare neutral. Notice

[11] This occurs, for example, when by charging a fixed fee $T_i > 0$ the monopolist gains T_i but the platform loses $(1 + \mu)T_i$, and μ is sufficiently large (see, e.g., Calzolari et al., 2018). See the Appendix for a complete analysis of this case.

also that while imposing a price-cap is equivalent to forcing prices (because intermediaries would like to set prices higher than M), imposing a price-floor is welfare neutral since it does not help M to solve the multiple marginalization problem.

RPM is welfare improving also when it is imposed by the platforms, and not by M. This is because platforms have an incentive to squeeze the intermediaries' mark-up, exactly as M would do. In this case platforms would still make positive profits because of double marginalization (see, e.g., also Gaudin, 2019). Therefore, following the same logic of our baseline model, platform parity would still be pro-competitive.[12]

In sum, in the wholesale model platform parity is likely to increase welfare also when retail price restrictions are imposed along the supply chain. We will see in the next section that this conclusion changes in the agency model, where the distribution of the bargaining power is different.

Price commitment. Up until now, we have assumed that M sets the price on the direct channel at the last stage of the game (i.e., simultaneously with the intermediaries). Would the results change when M can commit to p_d before the contracting takes place? In the Appendix we show that, when M can credibly commit to the price charged in the direct channel, it acts as a Stackelberg leader and sets that price efficiently regardless of whether platform parity is in place or not. Therefore, the beneficial effect of platform parity is the same as in the baseline model, except that it must be diluted from the channel externality which wipes out because of commitment.

Multiple Platforms. Finally, it should be clear that the logic of the pro-competitive mechanism described above does not change when M deals with $N > 2$ symmetric platforms. In fact, as seen for γ, it can be shown (see the Appendix) that the difference between the pass on rate with and without parity is increasing in N: as the market becomes more competitive, intermediaries are more sensitive to cost variations. As a result, the beneficial effect of platform parity is more pronounced as competition in the industry intensifies.

2.3 The agency model

We can now turn to analyze the *agency model*. In this business model a contract between M and P_i specifies a commission (fee) f_i paid by M to P_i for each unit distributed by I_i. In addition, M sets the access price τ_i that it charges I_i for every unit sold through

[12]The formal argument is standard and omitted for brevity. Proofs are available upon request.

P_i. Following the literature (e.g., Boik and Corts, 2016, and Johansen and Vergé, 2016, among others) the timing of the game is as follows (see Figure 2.4):

$t = 1$ Platforms secretly offer commissions to M;

$t = 2$ M accepts or refuses these offers, and sets the access prices;

$t = 3$ The monopolist and the intermediaries set final prices and demand is allocated across and within the two channels.

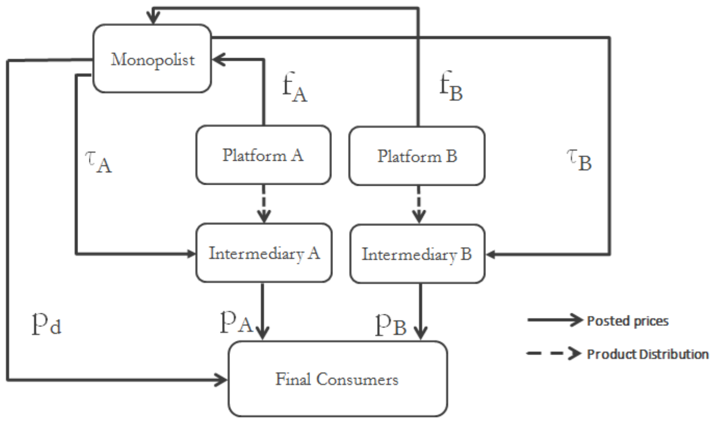

Fig. 2.4: The agency model.

We consider again two different versions of the game, depending on the possibility of signing a *platform parity agreement* between the monopolist and the platforms. Specifically:

- When a platform parity agreement is in place ($k = 1$), M commits to post the same access price on both platforms — i.e., $\tau_i = \tau$ for every $i = A, B$.
- Without the agreement ($k = 0$), M is free to charge intermediaries different access prices — i.e., it may happen that $\tau_A \neq \tau_B$.

As before, the equilibrium concept is *contract equilibrium* and we restrict our attention to symmetric equilibria in which, for every $k = 1, 0$:[13]

- Both platforms offer the same commission f_k^* to M;
- M accepts both offers and sets the same access price τ_k^* on each platform;

[13] As in the wholesale model, we focus directly on a symmetric equilibrium for simplicity of exposition and without loss of generality.

- Intermediaries charge the same retail price p_k^* to final consumers, while M charges $p_{d,k}^*$ on the direct channel.

It should be noted that, compared to the previous section, we have now changed both the direction of payment flows and the distribution of bargaining power — i.e., while in the wholesale model M had full bargaining power *vis-à-vis* platforms, now it is the opposite. The reason is intuitive: if M could set the commissions f_A and f_B, platforms would clearly make zero profit. The analysis would then be equivalent to the wholesale model: since platforms play no role, the parity provision would again unambiguously increase consumer surplus because it reduces the intermediaries' mark-ups. Hence, in giving full contractual power to the platforms against the monopolist, we are amplifying the anticompetitive effect of platform parity.

2.3.1 Equilibrium with platform parity

With platform parity each intermediary I_i solves

$$\max_{p_i} D^i\left(p_i, p_1^*, p_{d,1}^*\right)(p_i - \tau),$$

whose (standard) first-order condition is

$$\frac{1-\gamma-(1+\gamma)\,p_i+\gamma\left(p_1^*+p_{d,1}^*\right)}{(1-\gamma)(1+2\gamma)} - \frac{1+\gamma}{(1-\gamma)(1+2\gamma)}(p_i-\tau) = 0. \tag{2.7}$$

The solution of (2.7) yields the (symmetric) equilibrium of the (pricing) game between the intermediaries

$$p_1^*(\tau) \triangleq \frac{1+\gamma}{2+\gamma}\tau + \frac{1+\gamma(p_{d,1}^*-1)}{2+\gamma},$$

which, as expected, is increasing in the (common) wholesale price τ.
For any pair of commissions (f_A, f_B) negotiated with the platforms, M solves

$$\max_{p_d} \sum_{i=A,B} (\tau - f_i) D^i(p_1^*(\tau), p_1^*(\tau), p_d) + p_d D^d(p_d, p_1^*(\tau), p_1^*(\tau)),$$

whose first-order condition is

$$\underbrace{\frac{1-\gamma-(1+\gamma)\,p_d+2\gamma p_1^*\,(t)}{(1-\gamma)(1+2\gamma)} - p_d\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}}_{\text{Monopoly rule}} + \underbrace{\left(2\tau - \sum_{i=A,B} f_i\right)\frac{\gamma}{(1-\gamma)(1+2\gamma)}}_{\text{Channel externality}} = 0.$$

As in the wholesale model, this condition reflects both the standard monopoly trade-off and the channel externality due to the impact of p_d on the profit that M obtains through the indirect channel.

For any pair (f_A, f_B), let $p^*_{d,1}(\tau, f_A, f_B)$ denote the solution of the first-order condition above. Following a backward induction logic, M sets τ by solving

$$\max_{\tau} \sum_{i=A,B} (\tau - f_i) D^i(p^*_1(\tau), p^*_1(\tau), p^*_{d,1}(\tau, f_A, f_B)) + p^*_{d,1}(\tau, f_A, f_B) D^d(p^*_{d,1}(\tau, f_A, f_B), 2p^*_1(\tau)),$$

whose first-order condition (by the Envelope Theorem) is

$$\underbrace{2\frac{1-\gamma-p^*_1(\tau)+\gamma p^*_{d,1}(\tau, f_A, f_B)}{(1-\gamma)(1+2\gamma)}}_{\text{Revenue Enhancing}} - \underbrace{\sum_{i=A,B} (\tau - f_i) \frac{1}{(1-\gamma)(1+2\gamma)} \frac{\partial p^*_1(\tau)}{\partial \tau}}_{\text{Demand Reduction}} +$$

$$+ \underbrace{2p^*_{d,1}(\tau, f_A, f_B) \frac{\partial p^*_1(\tau)}{\partial \tau} \frac{\gamma}{1+\gamma(1-2\gamma)}}_{\text{Channel Externality}} = 0. \tag{2.8}$$

Hence, M's choice of τ is shaped by the following effects. First, a higher τ increases the retail prices on the indirect channel, which leads to a higher demand on the direct channel: a positive channel externality. Second, higher retail prices in the indirect channel reduce the demand faced by each intermediary: a demand reduction effect. Third, for a given demand, a higher τ increases the revenue that M collects from the intermediaries: a revenue enhancing effect.

Solving equation (2.8) yields M's optimal wholesale price $\tau^*_1(f_A, f_B)$. Thus, going back to the first stage of the game, we can solve each platform's maximization problem. Because the consumer displays preference for variety, we neglect for the moment M's participation constraint, which is verified in the Appendix. Hence, P_i solves

$$\max_{f_i} f_i D^i(p^*_1(f_i, f^*_1), p^*_1(f_i, f^*_1), p^*_{d,1}(f_i, f^*_1)),$$

where, to save on notation, we defined

$$p^*_{d,1}(\tau^*_1(f_i, f^*_1), f_i, f^*_1) \triangleq p^*_{d,1}(f_i, f^*_1),$$

and

$$p^*_1(\tau^*_1(f_i, f^*_1)) \triangleq p^*_1(f_i, f^*_1).$$

Differentiating with respect to f_i, we obtain

$$\underbrace{\frac{1-\gamma-p_1^*(f_i,f_1^*)+\gamma p_{d,1}^*(f_i,f_1^*)}{(1-\gamma)(1+2\gamma)}}_{\text{Revenue Enhancing}}-\underbrace{\frac{f_i}{(1-\gamma)(1+2\gamma)}\frac{\partial p_1^*(f_i,f_1^*)}{\partial f_i}}_{\text{Multiple Marginalization}}+$$

$$+\underbrace{\frac{\gamma f_i}{(1-\gamma)(1+2\gamma)}\frac{\partial p_{d,1}^*(f_i,f_1^*)}{\partial f_i}}_{\text{Business Stealing}}=0.$$

The interpretation of this condition is as follows. First, holding demand constant, when P_i charges M a higher fee, it earns higher revenues. Second, *ceteris paribus*, a higher f_i also induces M to charge a higher access price. Hence, retail prices in the indirect channel increase, whereby reducing the volume of sales on that platform: a multiple mark-ups problem. Third, since a higher f_i exacerbates the marginalization problem, M has an incentive to reduce the price on the direct channel in order to increase demand on the (relatively) more efficient direct channel, whereby reducing P_i's demand.
Imposing symmetry, we can state the following.

Proposition 2.4. *With platform parity, the symmetric equilibrium of the agency model has the following features:*

(*i*) *Each platform charges* M

$$f_1^*=\underbrace{(1-\gamma)\frac{4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4}{6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4}}_{(>0)\ P_i\text{'s mark-up}}.$$

(*ii*) *The monopolist sets*

$$p_{d,1}^*=p^M+\underbrace{\frac{\gamma(1-\gamma^2)}{2(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4)}}_{\text{Channel externality}},$$

and charges each intermediary

$$\tau_1^*=f_1^*+\underbrace{\frac{2+9\gamma+11\gamma^2-6\gamma^3-13\gamma^4-4\gamma^5}{2\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)}}_{(>0)\ M\text{'s mark-up}},$$

with $\tau_1^*>f_1^*$.

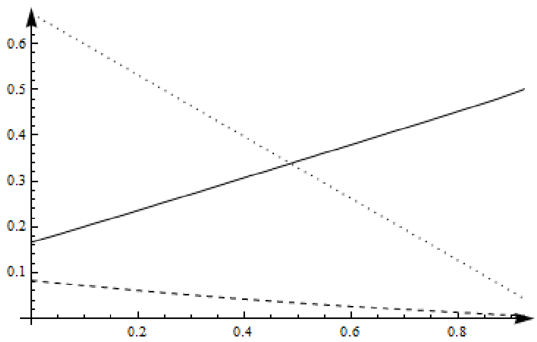

Fig. 2.5: Intermediaries' (dashed line), platforms' (dotted line) and M's (continuous line) markups as functions of γ in the agency model with a platform parity agreement.

(iii) The intermediaries set

$$p_1^* = \tau_1^* + \underbrace{\frac{(1-\gamma)^2(1+\gamma)(1+2\gamma)}{2\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)}}_{(>0)\ I_i\text{'s mark-up}},$$

with $p_1^* \geq p_{d,1}^*$ *and* $p_1^* \geq \tau_1^*$.

The equilibrium of the game features multiple mark-ups even in the agency model. The reason is straightforward: in order to earn profits platforms must charge positive commissions to M, who is thus forced to pass on these higher commissions to the intermediaries. As a result, in the equilibrium, intermediaries set a retail price that is higher than the price charged by the multi-product monopolist. Moreover, M will also charge a higher price on the direct channel because demand in that segment increases in response to higher prices in the indirect channel.
Figure 2.5 shows that while the intermediaries and the platforms' mark-ups are decreasing in γ, as implied by tougher competition between and within the two distribution channels, M's mark-up is increasing in γ because more competition reduces the multiple mark-ups problem, thus M behaves more efficiently.

2.3.2 Equilibrium without platform parity

Next, assume that there is no platform parity agreement — i.e., M can charge different access prices to the intermediaries. Recall that contracts are secret: I_i observes only τ_i but not the access price charged to its rival. Hence, taking as given the equilibrium price of its competitors — i.e., p_0^* and $p_{d,0}^*$ — I_i solves

$$\max_{p_i} D^i\left(p_i, p_0^*, p_{d,0}^*\right)(p_i - \tau_i),$$

whose first-order condition immediately yields

$$p_0^*(\tau_i) \triangleq \frac{\tau_i}{2} + \frac{1+\gamma p_0^* - \gamma(1-p_{d,0}^*)}{2(1+\gamma)}.$$

As intuition suggests, $p_0^*(\tau_i)$ is increasing in τ_i and (due to strategic complementarity) in p_0^* and $p_{d,0}^*$.

In stage 3, after observing f_A and f_B negotiated with the platforms, M sets the price p_d in order to solve

$$\max_{p_d} \sum_{i=A,B} (\tau_i - f_i) D^i(p_0^*(\tau_i), p_0^*(\tau_{-i}), p_d) + p_d D^d(p_d, p_0^*(\tau_A), p_0^*(\tau_B)).$$

The first-order condition is

$$\underbrace{\frac{1-\gamma-(1+\gamma)\,p_d+\gamma\,(p_0^*(\tau_A)+p_0^*(\tau_B))}{(1-\gamma)(1+2\gamma)} - p_d\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}}_{\text{Monopoly rule}}+$$

$$+\underbrace{\frac{\gamma}{(1-\gamma)(1+2\gamma)}\sum_{i=A,B}(\tau_i - f_i)}_{\text{Channel externality}} = 0.$$

Once again, this condition reflects the standard monopoly trade-off and the channel externality.

Let $p_{d,0}^*(\tau_A, \tau_B, f_A, f_B)$ denote the solution to the considered first-order condition. At stage 2, M chooses τ_A and τ_B in order to solve

$$\max_{\tau_A,\tau_B} \sum_{i=A,B} (\tau_i - f_i) D^i(p_0^*(\tau_i), p_0^*(\tau_{-i}), p_{d,0}^*(\cdot)) + p_{d,0}^*(\cdot) D^d(p_{d,0}^*(\cdot), p_0^*(\tau_A), p_0^*(\tau_B)).$$

By the Envelope Theorem, the first-order condition with respect to τ_i ($i = A, B$) is

$$\underbrace{\frac{1-\gamma-(1+\gamma)\,p_0^*(\tau_i)+\gamma\left(p_0^*(\tau_{-i})+p_{d,0}^*(\cdot)\right)}{(1-\gamma)(1+2\gamma)} - \frac{1+\gamma}{(1-\gamma)(1+2\gamma)}(\tau_i - f_i)\frac{\partial p_0^*(\tau_i)}{\partial \tau_i}}_{\text{Monopoly rule}}+$$

$$+\underbrace{\frac{\gamma}{(1-\gamma)(1+2\gamma)}\left[\tau_{-i} - f_{-i} + p_{d,0}^*(\cdot)\right]\frac{\partial p_0^*(\tau_i)}{\partial \tau_i}}_{\text{Channel externality}} = 0.$$

Clearly, M's profit increases as τ_i grows large because (for a given demand) M collects higher revenues from I_i. However, since contracts are secret, a higher τ_i only increases the retail price charged by I_i. Hence, other things being equal, I_i's demand drops, whereas I_{-i}'s demand on the indirect channel and M's demand on the direct channel increase. By solving the system of first-order conditions, we obtain the monopolist's price on platform P_i as a function $\tau_0^*(f_i, f_{-i})$. To save on notation, let

$$p_{d,0}^*(\tau_0^*(f_i, f_{-i}), \tau_0^*(f_{-i}, f_i), f_i, f_{-i}) \triangleq p_{d,0}^*(f_i, f_{-i}),$$

and

$$p_0^*(\tau_0^*(f_i, f_{-i})) \triangleq p_0^*(f_i, f_{-i}).$$

Then, P_i solves the following maximization problem at stage 1:

$$\max_{f_i} f_i D^i(p_0^*(f_i, f_0^*), p_0^*(f_0^*, f_i), p_{d,0}^*(f_i, f_0^*)),$$

whose first-order condition is

$$\underbrace{\frac{1-\gamma-(1+\gamma)\,p_0^*(f_i, f_0^*)+\gamma\left(p_0^*(f_0^*, f_i)+p_{d,0}^*(f_i, f_0^*)\right)}{(1-\gamma)(1+2\gamma)} - f_i\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}\frac{\partial p_0^*(f_i, f_0^*)}{\partial f_i}}_{\text{Monopoly rule}}+$$

$$\underbrace{+f_i\frac{\gamma}{(1-\gamma)(1+2\gamma)}\left[\frac{\partial p_0^*(f_0^*, f_i)}{\partial f_i}+\frac{\partial p_{d,0}^*(f_i, f_0^*)}{\partial f_i}\right]}_{\text{Channel externality}} = 0.$$

There are three effects that shape P_i's optimal commission. First, for a given demand, by increasing the commission charged to M, platform P_i earns a higher profit. Second, since (*ceteris paribus*) M reacts to a higher f_i by increasing the access price charged to I_i, fewer consumers demand P_i's product. Third, since the price charged by I_i increases, the demand on platform P_{-i} and on the direct channel increase. Intuitively, as I_i becomes less competitive, M has an incentive to reduce τ_{-i} and p_d.
Imposing symmetry we can state the following.

Proposition 2.5. *Without platform parity, the symmetric equilibrium of the agency model has the following features:*

(*i*) *Each platform charges M*

$$f_0^* = \underbrace{\frac{8\,(1-\gamma^2)\,(4+4\gamma-9\gamma^2)}{64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4}}_{(>0)\ P_i\text{'s mark-up}}.$$

(ii) The monopolist sets

$$p^*_{d,0} = p^M + \underbrace{\frac{\gamma(1+\gamma)(8-9\gamma^2)}{64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4}}_{(>0)\ \textit{Channel externality}},$$

and charges each intermediary

$$\tau^*_0 = f^*_0 + \underbrace{\frac{32+128\gamma+100\gamma^2-148\gamma^3-153\gamma^4}{2(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)}}_{(>0)\ M_i\textit{'s mark-up}},$$

with $\tau^*_0 > f^*_0$.

(iii) Each intermediary sets

$$p^*_0 = \tau^*_0 + \underbrace{\frac{(1-\gamma)(1+2\gamma)(8-9\gamma^2)}{64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4}}_{(>0)\ I_i\textit{'s mark-up}},$$

with $p^*_0 \geq p^*_{d,0}$ *and* $p^*_0 \geq \tau^*_0$.

Once again, the equilibrium features multiple marginalization. Platforms and intermediaries' mark-ups are decreasing in γ whereas M's mark-up is increasing in γ (see Figure 2.6).

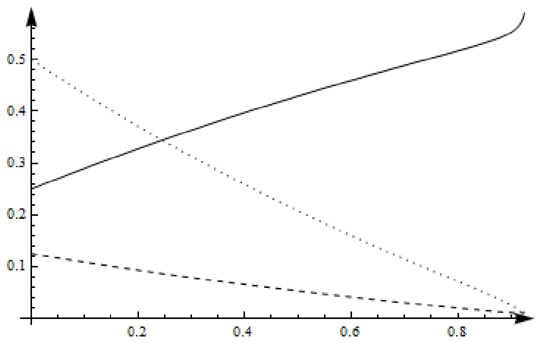

Fig. 2.6: Intermediaries' (dashed line), platforms' (dotted line) and M's (continuous line) markups as functions of γ in the model without platform parity agreements.

2.3.3 Welfare

We can now study how the introduction of a platform parity agreement affects consumer surplus and industry profits. To begin with, it is useful to examine the effect of such an agreement on commissions, access prices and final (retail) prices.

Proposition 2.6. *With platform parity:*

(i) *P_i $(i = A, B)$ charges M a higher commission — i.e., $f_1^* > f_0^*$;*

(ii) *M always charges a lower price on the direct channel, but it charges a higher access price if and only if γ is not too large — i.e., $p_{d,1}^* < p_{d,0}^*$ for every $\gamma \in [0, \overline{\gamma}]$, while there is $\check{\gamma} \in (0, \overline{\gamma})$ such that $\tau_1^* < \tau_0^*$ if and only if $\gamma \geq \check{\gamma}$;*

(iii) *I_i $(i = A, B)$ sets a higher (retail) price if and only if γ is not too large — i.e., there is $\gamma^* \in (0, \check{\gamma})$ such that $p_1^* < p_0^*$ if and only if $\gamma \geq \gamma^*$. Notice that $\gamma^* < \check{\gamma}$, so that $\tau_1^* < \tau_0^*$ implies $p_1^* < p_0^*$.*

The reason why commissions are higher under platform parity is as in Boik and Corts (2016): each platform anticipates that, being concerned with double marginalization, under the parity provision M has a lower incentive to pass on commissions to the intermediaries, because the increase in a platform's commission must be passed on to both intermediaries and not just to that platform's intermediary. Hence, they charge M more. Interestingly, the difference between the commissions with and without parity is decreasing in γ (see Figure 2.7). The reason is that, as products become more differentiated (lower γ), competition (within and across channels) intensifies and, as a result, mark-ups and profits in both regimes decrease, whereby aligning the equilibrium commissions with and without parity.

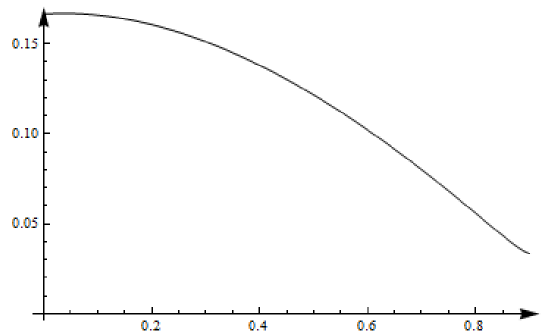

Fig. 2.7: Difference $f_1^* - f_0^*$ as a function of γ.

The effect of the provision on the price charged on the direct channel is straightforward. Since platform parity increases commissions, M charges a lower p_d in order to divert business towards the (relatively cheaper) direct channel. Notice, however, that diverting business towards direct distribution at the expense of direct distribution is not necessarily efficient for M because of consumers' taste for variety.

The effect of the provision on the access price is ambiguous. Since the parity provision induces platforms to charge higher commissions, M marks up more the intermediaries. But, by doing so, it also reduces their demand (since, in turn, intermediaries will charge higher retail prices) and thus the volume of indirect sales. Clearly, when γ is large enough, the second effect prevails, because with parity pass on rates are higher when γ is large, whereas the opposite is true without parity.
Finally, the effect of platform parity on the retail prices is ambiguous too since it reflects the non-monotone impact of γ on the access price. Of course, whenever the provision lowers the equilibrium access price, it also lowers the retail price in the indirect channel (since it implies lower marginal costs for the intermediaries).
Summing up, in the agency model there are three main welfare effects associated with platform parity. First, other things being equal, by increasing commissions, the parity provision harms consumers because it creates more marginalization: the dark side of platform parity. Second, other things being equal, as in the wholesale model,s the provision mitigates the marginalization problem between the monopolist and the intermediaries, which benefits consumers: the bright side of platform parity. Third, by lowering prices on the direct channel, the provision increases consumer surplus in that segment and creates a competitive pressure on the indirect channel, which (*ceteris paribus*) benefits the consumers purchasing from the intermediaries.
Building on these insights, we can state the following.

Proposition 2.7. *There are two thresholds $\tilde{\gamma} \in (0, 1)$ and $\hat{\gamma} \in (0, 1)$, with $\hat{\gamma} > \tilde{\gamma}$, such that the introduction of platform parity:*

(*i*) *benefits consumers if and only if $\gamma \geq \tilde{\gamma}$;*

(*ii*) *always damages the monopolist and the intermediaries;*

(*iii*) *benefits platforms if and only if $\gamma \geq \hat{\gamma}$.*

Hence, platform parity benefits consumers for γ large. The intuition is the following. Intensified competition increases the difference between τ_0^* and τ_1^* and, at the same time, it reduces the difference between f_1^* and f_0^*. In other words, as in the wholesale model, the wedge between the rate at which intermediaries pass on with and without parity is increasing in γ — i.e.,

$$\frac{\partial p_1^*(\tau)}{\partial \tau} - \frac{\partial p_0^*(\tau_i)}{\partial \tau_i} = \frac{\gamma}{2\,(2+\gamma)},$$

whereas the wedge between the commissions decreases (as shown in Figure 2.7). Moreover, intensified competition also lowers prices in the direct channel, which benefits consumers on both channels.

Interestingly, in contrast with the wholesale benchmark, in the agency model platforms benefit from the parity provision only if γ is sufficiently large. The reason is that, when γ is large enough, platform parity exerts a strong downward pressure on the mark-ups, whereby allowing platforms to increase commissions without producing an effect that is too negative on demand. The same argument explains why the monopolist and the intermediaries are hurt by the provision. The monopolist gains lower margins and pays higher commissions. The intermediaries earn lower margins because prices on the direct channel are lower with rather than without the provision and (for γ large) prices on the indirect channel are lower with the provision.

Hence, a novel policy implication of our analysis is that platforms' incentives are aligned with consumers' interests since — i.e., whenever the parity provision benefits platforms, it also benefits consumers ($\hat{\gamma} > \tilde{\gamma}$), but not the other way around. Hence, from a practical point of view, the likelihood that these provisions benefit consumers is higher when their introduction is not opposed by platforms. By contrast, the monopolist and the intermediaries's incentives are misaligned with consumer surplus.

Figure 2.8 describes the impact of platform parity on profits and consumer surplus.

Finally, in the next proposition we examine to what extent platforms can use side payments to align their incentives concerning the parity regime with the other players in the industry.

Proposition 2.8. *There exist thresholds $\gamma_T > \gamma_M > \gamma_I > \hat{\gamma}$ such that the introduction of platform parity:*

- *increases the joint profit of platforms and intermediaries if and only if $\gamma > \gamma_I$;*
- *increases the joint profit of platforms and the monopolist if and only if $\gamma > \gamma_M$;*
- *increases the total industry profit if and only if $\gamma > \gamma_T$.*

The intuition is straightforward. When competition in the product market is sufficiently fierce, the beneficial effect of the provision on the profit of the platforms outweighs the negative effect on the profit of the intermediaries. In fact, for γ large, the provision not only allows platforms to increase commissions without producing a too negative effect on demand, but it also reduces the access price charged by M to the intermediaries — i.e., $\tau_1^* < \tau_0^*$. Hence, platforms can persuade intermediaries to accept the parity regime by making appropriate side payments that align their incentives. However, persuading the monopolist is harder: whenever the joint profit of the platforms and the intermediaries

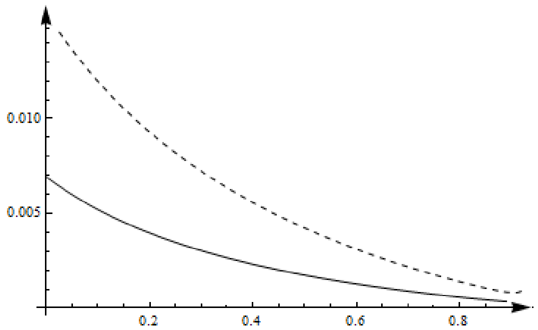

(a) *Intermediaries' profit in the presence of a platform parity clause (continuous line) and without parity agreements (dashed line).*

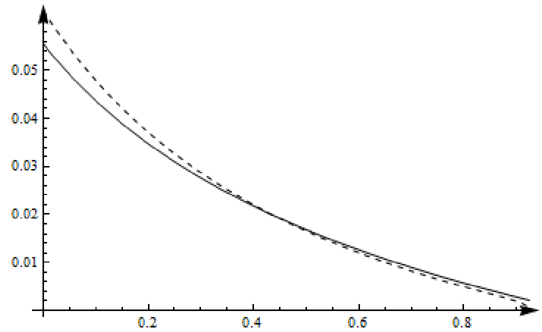

(b) *Platforms' profit in the presence of a platform parity clause (continuous line) and without parity agreements (dashed line).*

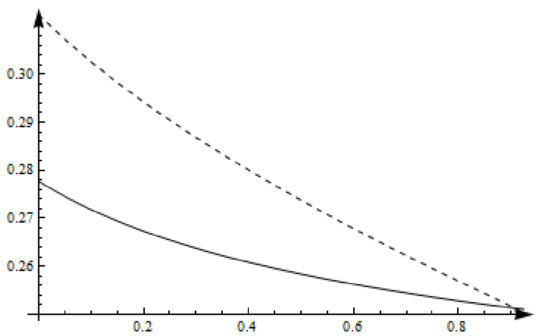

(c) *Monopolist's profit in the presence of a platform parity clause (continuous line) and without parity agreements (dashed line).*

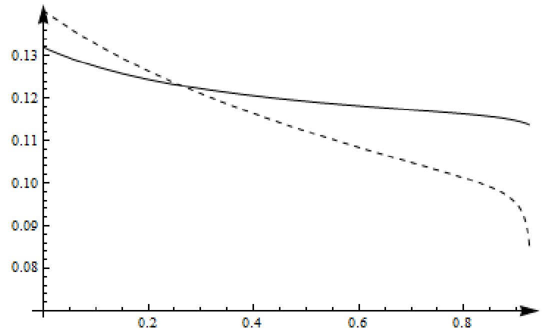

(d) *Consumer surplus in the presence of a platform parity clause (continuous line) and without parity agreements (dashed line).*

Fig. 2.8: Equilibrium profits and consumer surplus in the agency model as functions of γ.

is higher with than without the parity provision, the total industry profit may well be lower without the provision — i.e., the price that M would require to accept the parity regime is higher than the gain obtained jointly by the platforms and the intermediaries. As a result, in these cases, the only way to increase consumer surplus is to allocate more decision rights to the platforms than the monopolist, who is the most resilient player in the industry to the introduction of a consumer welfare enhancing parity provision.

2.3.4 Wholesale *vs* agency model

In this section we compare firms' profits and consumer surplus in the two business models examined before. We consider two alternative scenarios. First, we suppose that, in both business models, platforms (cooperatively) choose whether or not requiring platform parity. In this case, from the foregoing analysis, it follows that platforms would always choose to require the provision if the wholesale model is adopted, whereas in the agency model they would do so only for $\gamma \geq \hat{\gamma}$. Hence, to obtain a meaningful comparison between the two business models, in the following, we compare

- firms' profits and consumer surplus with platform parity in the wholesale model and without platform parity in the agency model in the region of parameters where $\gamma < \hat{\gamma}$;
- firms' profits and consumer surplus under platform parity in both business models in the region of parameters where $\gamma \geq \hat{\gamma}$.

Alternatively, we suppose that the monopolist, instead of platforms, decides whether to introduce platform parity, within each business model. In this case, from the foregoing analysis, we know that, for every $\gamma \in [0, \overline{\gamma}]$, M would introduce platform parity in the wholesale model and would not do so in the agency model.
The comparison between the two business models is summarized by the following Proposition.

Proposition 2.9. *Suppose that, within each business model, the presence of platform parity is dictated either by M or cooperatively by the platforms. In both cases, the two business models compare as follows:*

- *M is better off in the wholesale model, unless γ is sufficiently large;*
- *platforms are better off in the agency model;*
- *intermediaries are better off in the wholesale model;*
- *consumers are better off in the agency model.*

The result concerning consumer surplus reflects the fact that in the wholesale model the multiple marginalization problem is more pronounced than in the agency model. Intuitively, in the wholesale model platforms can influence the intermediaries' mark-ups through the access price, and when setting this price they only internalize the provision's effect on their own demand, and not the externality on the direct channel. By contrast, in the agency model M directly controls the intermediaries' mark-ups (by choosing the access price) and has a clear incentive to reduce multiple marginalization to minimize the (negative) channel externality. As a consequence, as in Johnson (2017), regardless of the parity regime, final prices are lower in the agency model than in the wholesale model.

Not surprisingly, and in line with the results of Johnson (2017), platforms gain more in the agency model, since in this case they exploit a first-mover advantage *vis-à-vis* the monopolist. Interestingly, however, M itself can be better off in the agency model. The trade-off is as follows. As pointed out above, regardless of the parity regime, shifting from the wholesale to the agency model reduces the multiple marginalization problem, which, *ceteris paribus*, benefits M; however, in the agency model M loses bargaining power *vis-à-vis* platforms. Since in the agency model (in both parity regimes), platforms' commissions are decreasing in γ, it follows that, overall, M is worse off in that business model when γ is relatively small, and better off otherwise. Moreover, since, in the agency model, platforms' commissions are higher with platform parity, the range of γ where the adoption of the agency model benefits M enlarges when the provision is not in place — i.e., when M can choose whether to introduce platform parity. Clearly, intermediaries prefer the wholesale model because final prices in the agency model are lower: with an agency structure (in equilibrium) there is more competition both within and across distribution channels.

Interestingly, platforms' incentives over the choice of the business model are aligned with those of consumers. Specifically, if platforms can (cooperatively) choose the business model, then, regardless of which player is then allowed to dictate the parity regime, platforms would always choose the agency model, which maximizes consumer surplus.

2.3.5 Commitment

In this section we examine the agency model under the assumption that M can credibly commit to the price on the direct channel before contracting with platforms. Clearly, the intermediaries' maximization problem is the same as in the no commitment scenario analyzed above, except for the fact that they now observe p_d before setting their price. The analysis follows the same backward logic as in the wholesale model. Hence, for brevity, we relegate the details to the Appendix and state here only the main result.

Proposition 2.10. *If M can credibly commit to the price on the direct channel before contracting with platforms, it chooses $p^*_{d,k} = p^M$ for $k = 1, 0$. Platforms charge higher commissions under platform parity and final prices on the indirect channel are still upward distorted with respect to p^M — i.e.,*

$$p_1^* = p^M + \underbrace{\frac{(1-\gamma)(5+2\gamma)}{6(2+\gamma)}}_{\textit{Multiple marginalization}} > p_0^* = p^M + \underbrace{\frac{(1-\gamma)(6+\gamma)}{2(8+3\gamma)}}_{\textit{Multiple marginalization}}.$$

Hence, with commitment, platform parity always harms consumers in the agency model.

When M commits to the efficient price p^M, the beneficial effect of the parity provision is completely dissipated. The reason is the following. The introduction of a price parity provision induces the platforms to increase the commissions charged to the monopolist, who is thus forced to increase the access price charged to the intermediaries in order to mark up these higher commissions. However, under commitment, M cannot compensate the effect of increased marginalization with a lower price in the direct channel because that price is set efficiently regardless of whether the provision is in place or not. Hence, with price commitment, platform parity always harms consumers.

2.3.6 $N > 2$ competing platforms

We now investigate the effect of increased competition in the indirect channel in the agency model. To do so, we assume that M deals with $N \geq 2$ symmetric platforms, each being in an exclusive relationship with one intermediary. Following the previous approach, consider a representative consumer whose preferences are described by the following linear-quadratic utility function

$$U(\cdot) \triangleq \sum_{i\in\mathcal{I}} q_i - \frac{1}{2}\sum_{i\in\mathcal{I}} q_i^2 - \gamma \sum_{i,j\in\mathcal{I},j\neq i} q_i q_j - \sum_{i\in\mathcal{I}} p_j q_j + m,$$

where $\mathcal{I} \triangleq \{1, \ldots, N, d\}$. Standard techniques then yield the (direct) demand functions

$$q_i(\cdot) = \frac{1-\gamma}{1+(N-1)\gamma - N\gamma^2}\left(1 - \frac{1+(N-1)\gamma}{1-\gamma}p_i + \frac{\gamma}{1-\gamma}\sum_{j\in\mathcal{I}\setminus\{i\}} p_j\right).$$

Finally, as before, we assume that platform parity either applies to all platforms or none. For brevity we only provide a graphical analysis of the results that are detailed in the Appendix. To begin with, it is useful to understand how N affects the bright side of platform parity — i.e., how the wedge between the rate at which intermediaries pass on with and without parity varies with N. Figure 2.9 panel (a) shows that, for a given γ, as the number of platforms increases the bright side of platform parity amplifies. The intuition is as before: increased competition (as reflected by a larger N) leads intermediaries to be more responsive to 'marginal costs', which refrains M from charging higher access prices. However, Figure 2.9 panel (b) shows that also the dark side of the provision increases with N — i.e., the difference between $f_1^* - f_0^*$ is positive and increasing in N. The reason is again related to the impact of the parity provision on the multiple marginalization effect: other things being equal, as N grows large the rate at which intermediaries pass

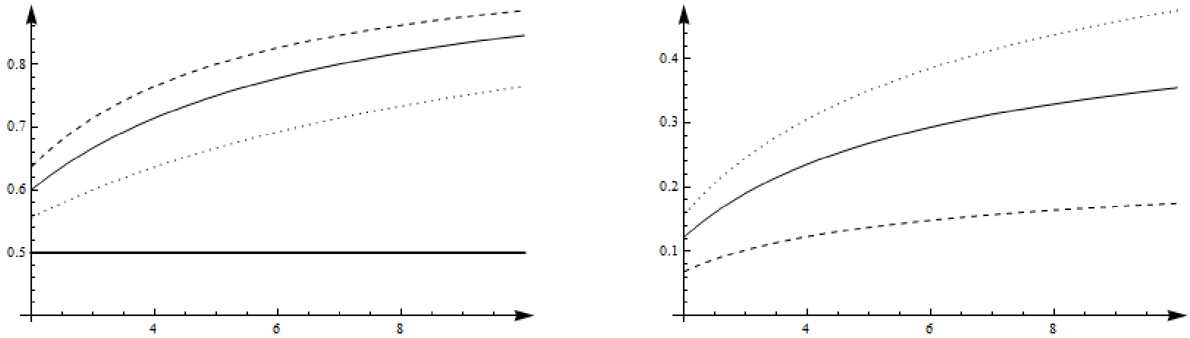

(a) *Derivatives* $\frac{\partial p_0^*(\tau_i)}{\partial \tau_i}$ *(thick line) and* $\frac{\partial p_1^*(\tau)}{\partial \tau}$ *as functions of* $N \geq 2$.

(b) *Difference* $f_1^* - f_0^*$ *as function of* $N \geq 2$.

Fig. 2.9: Parameter's value: $\gamma = .25$ (dotted lines), $\gamma = .5$ (continuous lines), $\gamma = .75$ (dashed lines).

on costs is higher with than without platform parity. This greater responsiveness implies that commissions are passed on by the monopolist to the intermediaries at a rate that is (*ceteris paribus*) lower with platform parity and, even more so, when N is large. As a result, the difference between the commissions charged by the platforms to the monopolist with and without the provision is increasing with N.

Figure 2.10 illustrates the net effect of an increase of N on consumer surplus. It can be seen that for sufficiently large values of γ, the positive effect of N on the bright side of the provision dominates the effect on the dark side. By contrast, for values of γ below $\tilde{\gamma}$ — i.e., when the parity provision damages consumers in duopoly — a higher N tends to make the negative effect of the provision on consumer surplus even worst.

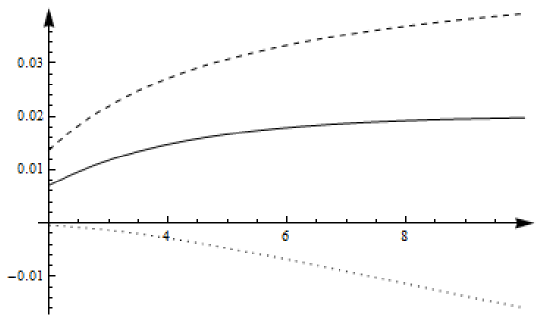

Fig. 2.10: Difference between the consumer surplus in the presence and in the absence of a platform parity agreement as a function of $N \geq 2$. Parameter's value: $\gamma = .25$ (dotted lines), $\gamma = .5$ (continuous lines), $\gamma = .75$ (dashed lines).

2.3.7 Quality deterioration as a threat to break down parity

In this section we consider the case in which the products distributed in the indirect channel are of lower quality than the product sold on the direct channel — i.e., we assume that the two channels are vertically differentiated. The goal is to understand whether M has an incentive to distort quality in the indirect channel in order to induce platforms to drop the parity regime.

To model this feature in the simplest possible way, consider the following utility function of the representative consumer:

$$U(\cdot) \triangleq q_d + (1-c) \sum_{j=A,B} q_j - \frac{1}{2} \sum_{j=A,B,d} q_j^2 - \gamma \sum_{i,j=A,B,d;,j\neq i} q_j q_i - \sum_{j=A,B,d} p_j q_j + m,$$

where the parameter $c \in [0, 1-\gamma]$ can be interpreted as a measure of *quality deterioration*: with $c = 0$ there is no vertical differentiation between the two channels, but as c increases (*ceteris paribus*) consumers' utility from purchasing on the indirect channel drops.[14] For simplicity, we assume that the value of c is common knowledge. Standard techniques then yield the (direct) demand functions

$$q_i = \frac{1-c-\gamma-(1+\gamma)\,p_i + \gamma\,(p_{-i}+p_d)}{(1-\gamma)(1+2\gamma)}, \quad \forall i = A, B$$

$$q_d = \frac{1-\gamma+2c\gamma-(1+\gamma)\,p_d+\gamma\,(p_A+p_B)}{(1-\gamma)\,(1+2\gamma)}.$$

Hence, a lower quality (as reflected by a higher c) induces a parallel downward shift of the demand functions for the products distributed in the indirect channel, and an upward shift of the demand for the product directly distributed by M. Therefore, if we consider c as an exogenous parameter of the model, the analysis is the same as in the baseline model — i.e., the first-order conditions of the players' maximization problems have the same structure as before.

Let $\pi_k^M(c)$ and $\pi_k^P(c)$ denote M and the platforms' equilibrium profits for a given parity regime $k = 0, 1$, respectively. We can state the following.

Proposition 2.11. *$\pi_k^M(c)$ and $\pi_k^P(c)$ are decreasing in c for every $k = 0, 1$. Moreover,*

(i) $\pi_0^M(c) > \pi_1^M(c)$ *for every* $c \leq 1-\gamma$;

(ii) $\pi_0^P(c) \geq \pi_1^P(c)$ *if and only* $\gamma \leq \hat{\gamma}$.

[14] The restriction $c \leq 1-\gamma$ guarantees that equilibrium quantities are positive.

The intuition of this result is straightforward. When the quality of the product distributed in the indirect channel decreases, consumers are less willing to pay for that product: demand in the indirect channel drops. As a result, the monopolist and the platforms are worse off. Moreover, since a lower quality only shifts demand down but does not affect its slope, it does not influence the rate at which firms pass on costs and thus the extent of multiple marginalization. Hence, the results of the baseline model apply: M always prefers the regime without parity, whereas platforms prefer the provision as long as competition in the product market is sufficiently fierce.

These results suggest that, while M would always be better off by dropping parity, platforms may still want it for γ sufficiently large. However, when M can choose quality — i.e., when c is endogenous — it may use this choice as a strategic variable to align its incentives with the platforms'. Specifically, imagine that at the outset of the game M can commit to a menu of qualities depending on the parity regime which is then chosen cooperatively by the platforms — i.e., a menu (c_0, c_1) that specifies the quality of the product supplied on the indirect channel in every regime $k = 0, 1$.[15] After this preliminary stage, the game evolves as in the baseline model. We can state the following.

Proposition 2.12. *For every $\gamma > \hat{\gamma}$, there is a threshold $\bar{c} \in (0, 1-\gamma)$ such that, by choosing $c_0 = 0$ and any $c_1 = c > \bar{c}$, M induces platforms to drop parity — i.e., $\pi_0^P(0) > \pi_1^P(c)$.*

Therefore, if M can commit to a menu of qualities as a function of the parity regime, it can induce the worse choice for consumers.[16] Clearly, this result holds only if M has enough commitment power *vis-à-vis* the platforms. The reason is that, if M lacks such commitment power, choosing $c_1 = c > \bar{c}$ is not subgame perfect: if platforms choose to introduce parity, M would then have an incentive to set $c_1 = 0$ if this choice can be reneged. In practice, however, even if the monopolist lacks commitment power, repeated interactions can allow it to build a reputation that makes the threat of reducing quality in the parity regime credible. We plan to examine the interesting dynamic aspects of platform parity in future works.

2.3.8 Further remarks

Alternative bargaining. In the previous model we assumed that platforms make offers to the monopolist. The analysis considerably simplifies in the opposite scenario where M

[15] Clearly, as explained before, if M could choose both quality and parity regime, it would choose $c = 0$ and $k = 0$.

[16] Clearly, in the wholesale model quality is not an issue because players' incentives are aligned — i.e., the monopolist, the platforms and consumers are better off under the parity provision.

has full bargaining power and proposes contracts to the platforms. As intuition suggests, in this case M sets $f_i = 0$ in the equilibrium and platforms make zero profit. Hence, the main source of inefficiency hinges on the intermediaries' mark-ups. The logic is the same as in the wholesale model: M must charge a positive access price in order to make positive profits, which in turn induces the intermediaries to set excessively high retail prices. Platform parity is then a commitment device to mitigate double marginalization. Noteworthy, the same outcome realizes if M bilaterally bargains with each platform P_i over the commission f_i that maximizes their joint profit, and the bilateral surplus is then split by means of a fixed payment between M and P_i. Details are in the Appendix.

RPM and Full Content Agreements. What would happen if M can fix retail prices in the agency model? As argued in the wholesale model, RPM (or even a price-cap) also in the agency model allows M to eliminate the intermediaries' mark-ups (e.g., by setting $\tau_{i,k} = p_{i,k}$ for every $i = A, B$ and $k = 0, 1$). In such a model, however, as in Johansen and Vergé (2016), wide and narrow price parity agreements are anti-competitive. By a similar logic, platform parity on the top of RPM cannot increase consumer welfare (in contrast to the wholesale model where platform parity is always welfare neutral).[17]

By contrast, the logic of our model suggests that full content agreements — i.e., agreements according to which the monopolist negotiates with the platforms a retail price to be charged in the direct distribution channel in exchange of lower fees — are likely to be efficient. In fact, while those agreements may limit the moderating effect on platform's fees of competition from the direct distribution channel, the monopolist will not give up the leverage provided by its direct distribution channels unless platforms reduce their fees by as much as they would do in the absence of a full content agreement. Hence, access and retail prices need not be higher under a full content agreement. Furthermore, consumer welfare will tend to be higher if consumers can find the same content in their preferred distribution channel, which is precisely what a full content agreement ensures.

Interlocking relationships. As for the implications of our model in terms of antitrust policy, there is a further important reason why the *balance of harms* is likely to point in the direction of no intervention against platform parity provisions: unlike the model analyzed in this paper, platforms usually compete for the business of single-homing intermediaries as well as to increase their share of the business of multi-homing intermediaries. As a result, the presence of a platform parity agreement would incentivize platforms to lower their charges to the monopolist in order to gain market shares downstream. Hence, the likelihood that the net competitive effect of content parity is positive will be increased.

[17] The formal argument is standard and omitted for brevity. Proofs are available upon request.

Competing sellers. With competing sellers in the upstream market, the effect of platform parity produces different results depending on the business model under consideration. Specifically, while upstream competition unambiguously magnifies the welfare effects of the parity provision in the wholesale model (since it reduces wholesale prices), in the agency model the result depends on the degree of competition between direct and indirect distribution, as well as on the possibility of the sellers to develop their own direct distribution channels (in competition not only between them but also with the indirect distribution channels). In this respect, it is useful to recall that Johnson (2017) shows that when there is no direct channel, the same logic of Boik and Corts (2016) applies with competing sellers; actually, competition in the upstream market makes the anticompetitive effect of parity provisions even worst because platforms charge higher commissions to the sellers, who pass on these higher commissions to the intermediaries, whereby reducing consumer surplus. By contrast, Johansen and Vergé (2016) show that parity provisions may actually improve consumer surplus when competition in the upstream market is fierce enough and sellers, distributing their products also through their own direct channel, can delist from platforms charging commissions that are too high, and intensify competition through their direct channels.

Specific investments and the hold-up problem. We have also not considered the platforms' incentives to invest in activities that improve the quality of their services. The reason is again rather simple. Clearly, platforms have a greater incentive to increase quality as their profit grows large. Hence, in the wholesale model a parity provision increases welfare not only because it mitigates multiple marginalization, but also because it stimulates platforms' investments in quality. However, in the agency model the result depends on whether the provision increases or reduces platforms' profits, which as seen before depends on the extent of competition. The most interesting case is that in which $\gamma \geq \hat{\gamma}$, where the provision benefits both platforms and consumers. In this case, if the quality enhancing investment is a lump sum, M has a clear incentive to promise platform parity, wait for the investment to be undertaken, and then renege on the initial promise (of course, provided that the parity decision is in M's domain). Clearly, as long as M's renegotiation is anticipated by the platforms, the latter will not invest, whereby reducing consumer surplus.

Alternative solution concept. Following a recent literature, we have used Contract Equilibrium rather than PBE as a solution concept. The reason behind this choice is purely for tractability. As we have discussed before, Contract Equilibrium shares common features with passive beliefs, except for multilateral deviations which are often problematic for the existence of an equilibrium even with linear demand (see, e.g., Rey and Vergé, 2004). In addition, the three-layer structure of our model considerably complicates the analysis of PBE: in our game off-equilibrium beliefs should be specified not only in the contracting problem between the monopolist and the platforms (depending on the business model under consideration) but also for the intermediaries.

This makes our setting rather peculiar. To see why consider (for example) the wholesale model (the same logic applies *mutatis mutandis* to the agency model). When an intermediary (say I_i) observes an off-equilibrium contract, it faces the problem of conjecturing who has deviated along the supply chain. Specifically: does the unexpected contract reflect a deviation by M (and if so has I_{-i}'s contract been changed as well or not) or it simply reflects a deviation by P_i? The answer to this 'dilemma' has significantly different implications for the outcome of the game. In fact, if the deviation is attributed to P_i only, the 'no signal what you do not know principle' (Fudenberg and Tirole, 1991) implies that I_i has no reason to believe that I_{-i} has been offered a contract different than the equilibrium one: a logic coherent with passive beliefs. However, if the deviation is attributed to M, then I_i must believe that also I_{-i}'s contract has changed: a logic closer to the concept of wary beliefs (McAfee and Schwartz, 1994). Depending on which option one chooses, the equilibrium of the game may be rather complex to characterize either because wary beliefs are imposed twice in the game, or because they must coexist and be coherent with passive beliefs on the intermediaries' side. Nevertheless, the effect of platform parity seems robust to these issues: regardless of whether an equilibrium exists or not, with wary beliefs intermediaries are more responsive to unexpected offers than with passive beliefs because they are more suspicious about off-equilibrium contracts. This greater responsiveness tends to magnify the commitment power that M can gain under platform parity, and hence to preserve the qualitative insights of our analysis. We plan to address these interesting issues in the future (perhaps in a simpler environment).

General demand function. Although, for the sake of tractability, throughout the analysis we considered a linear demand function, the basic trade-off driving our results is robust to alternative specifications of the demand function. On the one hand, for every downward sloping demand function, platform parity yields higher platforms' commissions (see Boik and Corts, 2016) which, in turn, makes it more attractive for the monopolist to divert business towards the direct channel. On the other hand, also the bright side of platform parity is not specific to the assumption of linear demand. In fact, for every downward sloping demand function, with platform parity, an increase in the access price is equivalent to a common cost shock for the intermediaries, which is passed on to a greater extent as competition intensifies, whereas, without parity, a higher access price is equivalent to an idiosyncratic cost shock, which is passed on to a lesser extent when competition is fiercer.

2.4 Conclusion

In this paper, we have shown that in complex vertical industries — like, e.g., manufacturing industries with multi-tier supply chains (such as automotive, consumer appliances, electronic equipment, and apparel) or the airline ticket distribution industry — platform parity cannot be presumed anti-competitive in the absence of efficiencies.

We have seen that, if the wholesale model is adopted, as it is usually the case in traditional *brick-and-mortar* businesses, then these agreements are always pro-competitive, since they work as a commitment device to mitigate the multiple marginalization problem. Accordingly, not only consumers but also all firms in the supply chain benefit from these provisions. By contrast, when the agency model is employed, as it happens in many *two-sided markets*, even when platforms have full bargaining power in their relationship with the seller, being able to raise the commissions' level when a platform parity agreement is introduced, such a contractual provision might still be pro-competitive. In this case, consumers' preferences are always aligned with the platforms' but not with the seller's. Namely, as long as platforms benefit from platform parity agreements consumers gain as well. Hence, from a practical point of view, the likelihood that these provisions benefit consumers is higher when their introduction is not opposed by platforms.

In fact, we have also shown that while the platforms may persuade the intermediaries to accept the provision through appropriate side payments, persuading the monopolist is harder — i.e., even if the joint profit of the platforms and the intermediaries is higher with than without the provision, the total industry profit and the joint profit of the monopolist and the platforms may well drop when the provision is introduced. This implies that in the industry the incentives of the monopolist are the most likely not to be aligned with consumer surplus. Hence, in these cases, the only way to increase the welfare of the consumers is to allocate more decision rights on the industry structure to the platforms — i.e., to award them a stronger influence than the other players on whether the provisions under consideration should be banned or not.

Finally, we do acknowledge that in order to keep the analysis tractable we have sometimes neglected a few salient aspects of real life — e.g., upstream competition, merger incentives, quality provision with the related moral hazard problems, etc. We plan to explore these interesting issues in future research.

Appendix

Appendix 2.A. Proofs

Multi-product monopolist. The first-order conditions of the multi-product monopolist's problem with respect to p_i, $i = A, B$, and p_d are

$$\frac{1-\gamma-(1+\gamma)\,p_i+\gamma\,(p_{-i}+p_d)}{(1-\gamma)(1+2\gamma)}-\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}p_i+\frac{\gamma}{(1-\gamma)(1+2\gamma)}(p_{-i}+p_d)=0,$$

and

$$\frac{\gamma}{(1-\gamma)(1+2\gamma)}\sum_{i=A,B}p_i+\frac{1-\gamma-(1+\gamma)\,p_d+\gamma\sum_{i=A,B}p_i}{(1-\gamma)(1+2\gamma)}-\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}p_d=0,$$

respectively. Solving the system of these three first-order conditions, we obtain the price charged by the multi-product monopolist on the three channels. ■

Proof of Proposition 2.1. The result can be directly obtained by solving the first-order conditions derived in Section 2.2.1. ■

Proof of Proposition 2.2. The result can be directly obtained by solving the first-order conditions derived in Section 2.2.1. ■

Proof of Proposition 2.3. The comparison between the equilibrium prices in the two regimes is as follows:

$$p_1^*-p_0^*=-\frac{3\gamma(1+2\gamma\,(1+2\gamma))}{2(4+\gamma)(1+2\gamma)(8+5\gamma)}<0,$$

$$p_{d,1}^*-p_{d,0}^*=-\frac{3\gamma^2(5+2\gamma)}{2(4+\gamma)(1+2\gamma)(8+5\gamma)}<0.$$

In the presence of platform parity, equilibrium profits are

$$\pi_{i,1}^I=\frac{1-\gamma^2}{4(4+\gamma)^2(1+2\gamma)},$$

$$\pi_{i,1}^P=\frac{1-\gamma^2}{2(4+\gamma)^2(1+2\gamma)},$$

and

$$\pi_1^M=\frac{3(1+\gamma)(4+11\gamma+3\gamma^2)}{2(4+\gamma)^2(1+2\gamma)^2}.$$

In the game without platform parity, we have

$$\pi_{i,0}^{I} = \frac{1-\gamma^2}{(1+2\gamma)(8+5\gamma)^2},$$

$$\pi_{i,0}^{P} = \frac{2(1-\gamma^2)}{(1+2\gamma)(8+5\gamma)^2},$$

and

$$\pi_{0}^{M} = \frac{96+432\gamma+597\gamma^2+252\gamma^3}{4(1+2\gamma)^2(8+5\gamma)^2}.$$

Comparing these expression yields

$$\pi_{i,1}^{I} - \pi_{i,0}^{I} = \frac{3\gamma(16+7\gamma)(1-\gamma^2)}{4(4+\gamma)^2(1+2\gamma)(8+5\gamma)^2} > 0,$$

$$\pi_{i,1}^{P} - \pi_{i,0}^{P} = \frac{3\gamma(16+7\gamma)(1-\gamma^2)}{2(4+\gamma)^2(1+2\gamma)(8+5\gamma)^2} > 0,$$

$$\pi_{1}^{M} - \pi_{0}^{M} = \frac{9\gamma^2(8+98\gamma+103\gamma^2+22\gamma^3)}{4(4+\gamma)^2(1+2\gamma)^2(8+5\gamma)^2} > 0,$$

which concludes the proof. ■

Proof of Proposition 2.4. The equilibrium values are obtained by solving the first-order conditions derived in Section 2.3.1. However, to establish Propositions 2.4, it remains to be shown that the commissions values obtained by solving the first-order conditions of platforms' problem are such that M's participation constraint is satisfied. Specifically, since we assumed that M is allowed to delist from a platform, as in Johansen and Vergé (2016), the considered equilibrium exists if and only if, when the other platform P_{-i} offered the candidate equilibrium contract f_1^*, M does not find it profitable to decline P_i's offer f_1^*, thus being free to set prices on the two remaining channels — i.e., platform P_{-i} and the direct sale channel. If such a profitable deviation exists, then, in the symmetric contract equilibrium, platforms' commissions must be lower than f_1^*, and such that M's participation constraint is binding. Thus, in order to check whether a candidate equilibrium value f_1^* satisfies M's participation constraint, we should determine its profit when it declines P_i's offer, conditional on accepting the same offer f_1^* made by P_{-i}.

If M sells its product only through one platform, say P_i, then the linear-quadratic specification of the representative consumer's utility function becomes

$$U(q_i, q_d) \triangleq \sum_{j=i,d} q_j - \frac{1}{2}\sum_{j=i,d} q_j^2 - \gamma q_i q_d - \sum_{j=i,d} p_j q_j + m,$$

yielding the following demand functions

$$D(p_i, p_d)^i = \frac{1-\gamma(1-p_d)-p_i}{1-\gamma^2}, \quad D^d(p_d, p_i) = \frac{1-\gamma(1-p_i)-p_d}{1-\gamma^2}.$$

Assuming, for simplicity, that M's acceptance decisions are common knowledge, if M declined P_{-i}'s offer, then I_i solves[18]

$$\max_p D\left(p, \tilde{p}_d\right)\left(p-\tau\right),$$

where $\tilde{p}_d$ denotes the price that M is expected to set on the direct channel. Optimization yields a best-reply function

$$\tilde{p}(\tau) \triangleq \frac{\tau}{2} + \frac{1-\gamma(1-\tilde{p}_d)}{2}.$$

Accordingly, M sets τ and p_d by solving

$$\max_{\tau, p_d} p_d D^d(p_d, \tilde{p}(\tau)) + (\tau - f)D(\tilde{p}(\tau), p_d),$$

whose first-order conditions yield[19]

$$\tilde{\tau}(f) = \frac{4(1+f)-\gamma(4p_d^* + 3\gamma(1+\gamma+2f-\gamma\tilde{p}_d)-2)}{8-9\gamma^2},$$

and

$$\tilde{p}_d(f) = \frac{4-\gamma(f+\gamma(5+\tilde{p}_d)-1)}{8-9\gamma^2}.$$

Thus, to check whether M's participation constraint is satisfied when platforms set f_1^*, we must substitute this value into the above expressions (and impose $\tilde{p}_d(f_1^*) = \tilde{p}_d$). Let us first consider the case in which platform parity is in place. We obtain

$$\tilde{\tau}_1 = \frac{40+84\gamma-46\gamma^2-93\gamma^3-3\gamma^4+14\gamma^5}{8\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)},$$

and

$$\tilde{p}_d = p^M + \frac{\gamma(1-\gamma)(2+\gamma)(1+2\gamma)}{8(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4)},$$

from which it is easy to obtain M's *deviation profit*, which can finally be compared with the *equilibrium profit* — that is, the profit that M gets by accepting both platforms' offers. In Figure 2.11, we show that the equilibrium profit is, for every γ, higher than the

[18]Given that only one indirect channel is active, we omit subscripts, to ease notation.

[19]The corresponding second-order condition is satisfied for every $\gamma < \frac{2\sqrt{2}}{3}$.

deviation profit. ■

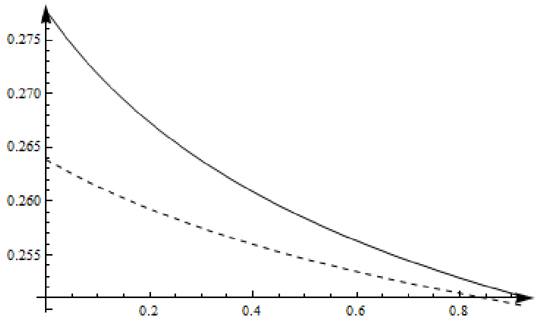

Fig. 2.11: M's equilibrium profit (continuous line) and deviation profit (dashed line) in the game with platform parity agreement as a function of γ.

Proof of Proposition 2.5. The equilibrium values are obtained by solving the first-order conditions derived in Section 2.3.2. It can be easily seen that M's participation constraint is not binding, in the game without platform parity, when platforms offer f_0^*.[20] ■

Proof of Proposition 2.6. The equilibrium values of commissions and prices in the two regimes, given in Propositions 2.4 and 2.5, are compared in Figure 2.12, which thus establishes Proposition 2.6. ■

Proof of Proposition 2.7. As for equilibrium profits, if a platform parity agreement is in place, we find

$$\pi_{i,1}^{I} = \frac{(1+2\gamma)\left(1-\gamma^2\right)^3}{4\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)^2},$$

$$\pi_{i,1}^{P} = \frac{\left(1-\gamma^2\right)^2\left(4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4\right)}{2\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)^2},$$

$$\pi_{i,1}^{M} = \frac{40+190\gamma+204\gamma^2-240\gamma^3-519\gamma^4-120\gamma^5+244\gamma^6+170\gamma^7+32\gamma^8}{4\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)^2},$$

whereas, without platform parity

$$\pi_{i,0}^{I} = \frac{(1-\gamma^2)(1+2\gamma)\left(8-9\gamma^2\right)^2}{(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)^2},$$

$$\pi_{i,0}^{P} = \frac{8(1-\gamma^2)(1+\gamma)\left(8-9\gamma^2\right)\left(4+4\gamma-9\gamma^2\right)}{(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)^2},$$

[20] As argued by Johansen and Vergé (2016), if a platform charges a higher commission, in the absence of platform parity, M finds it optimal to set a higher price on that platform, rather than shut down that channel. Therefore, in the regime without platform parity, the equilibrium value for the platforms' commission is the *unconstrained* one stated in Proposition 2.5.

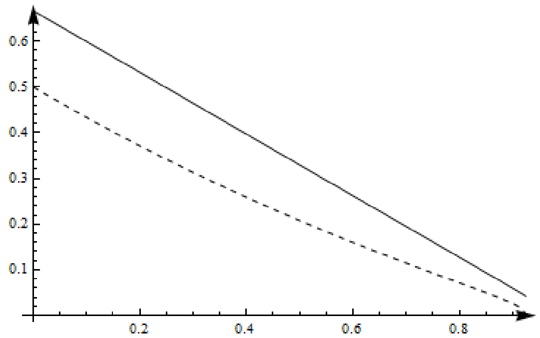

(a) *Equilibrium commissions* f_1^* *(continuous line) and* f_0^* *(dashed line) as functions of* γ.

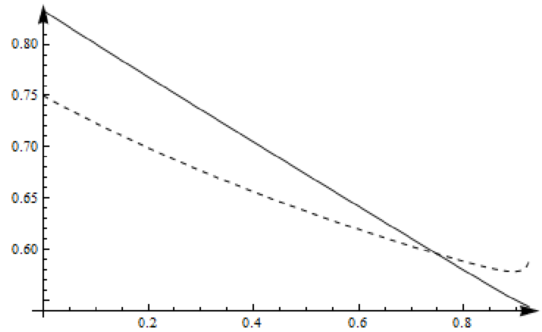

(b) *Equilibrium wholesale prices* τ_1^* *(continuous line) and* τ_0^* *(dashed line) as functions of* γ.

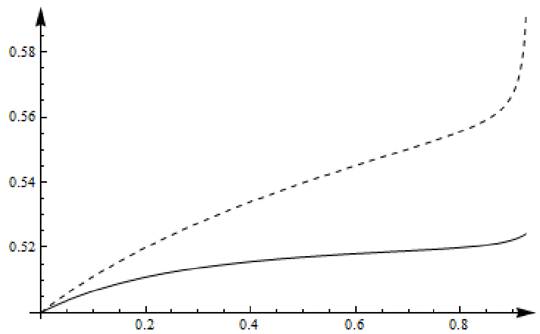

(c) *Equilibrium price* $p_{d,1}^*$ *(continuous line) and* $p_{d,0}^*$ *(dashed line) as functions of* γ.

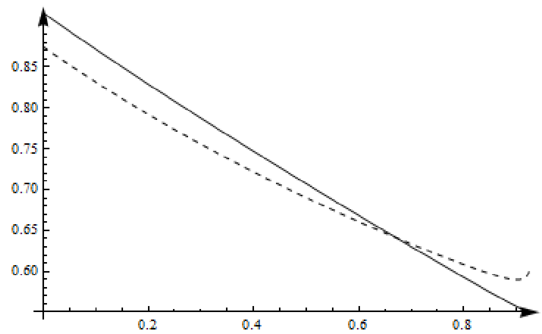

(d) *Equilibrium price* p_1^* *(continuous line) and* p_0^* *(dashed line) as functions of* γ.

Fig. 2.12: Equilibrium values as functions of γ.

$$\pi_{i,0}^M = \frac{1}{4} + \frac{(1+\gamma)^2 \left(8-9\gamma^2\right)^2 \left(4+4\gamma-9\gamma^2\right)}{(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)^2}.$$

These expressions are plotted in Figure 2.8, which establishes Proposition 2.7. ■

Proof of Proposition 2.8. The results follow from a direct comparison among the relevant joint profits in the two regimes: see Figure 2.13. ■

Proof of Proposition 2.9. The results easily follow from a direct comparison between the relevant profits. Details are available upon request. The comparison of consumer surplus is illustrated in Figure 2.14. ■

Proof of Proposition 2.10. In the game with commitment, the intermediaries observe p_d before setting their price. Therefore, in the presence of platform parity, their equilibrium price is a function of the (common) access price and of the price on M's direct channel, denoted by $p_1^*(\tau, p_d)$. Similarly, in the absence of platform parity, the best-reply function of I_i is denoted by $p_0^*(\tau_i, p_d)$.

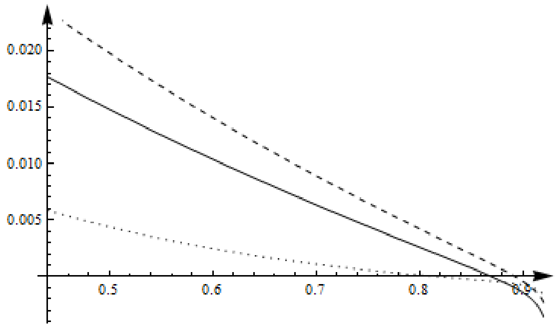

Fig. 2.13: Differences of joint profits $2[\pi^P_{i,0} + \pi^I_{i,0} - (\pi^P_{i,1} + \pi^I_{i,1})]$ (dotted line), $\pi^M_0 + 2\pi^P_{i,0} - (\pi^M_1 + 2\pi^P_{i,1})$ (continuous line), $\pi^M_0 + 2\pi^P_{i,0} + 2\pi^I_{i,0} - (\pi^M_1 + 2\pi^P_{i,1} + 2\pi^I_{i,1})$ (dashed line), as functions of $\gamma \in (\hat{\gamma}, \overline{\gamma})$.

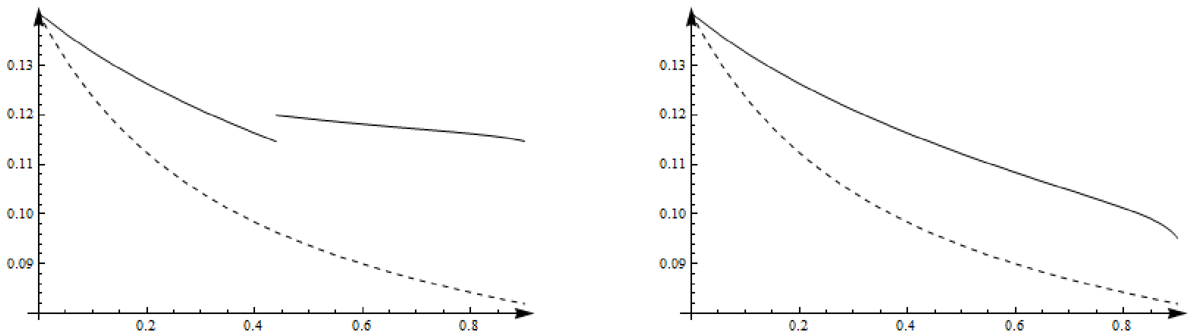

(a) *Parity regime chosen (cooperatively) by platforms. The discontinuity point in the agency model is at $\gamma = \hat{\gamma}$, in which the parity regime changes.*

(b) *Parity regime chosen by M.*

Fig. 2.14: Consumer surplus in the agency model (continuous line) and in the wholesale model (dashed line) as function of γ.

Equilibrium with platform parity. For any pair of commissions (f_A, f_B) negotiated with the platforms, and any announced value of p_d, M's optimization problem at the final pricing stage is as follows:

$$\max_{\tau} \sum_{i=A,B} (\tau - f_i) D^i(p^*_1(\tau, p_d), p^*_1(\tau, p_d), p_d) + p_d D^d(p_d, p^*_1(\tau, p_d), p^*_1(\tau, p_d)),$$

whose first-order condition is as in the baseline agency model of Section 2.3.1 and yields

$$\tau^*_1(f_i, f_{-i}, p_d) \triangleq \frac{2 + f_i + f_{-i} + 2\gamma(2p_d - 1)}{4}.$$

Moving backward at the previous stage of the game, P_i's problem is

$$\max_{f_i} f_i D^i(p_1^*(f_i, f_1^*, p_d), p_1^*(f_i, f_1^*, p_d), p_d),$$

where, to ease notation, $p_1^*(f_i, f_1^*, p_d) \triangleq p_1^*(\tau_1^*(f_i, f_1^*, p_d), p_d)$. The first-order condition for this optimization problem (which, again, is as in the corresponding problem without commitment), imposing symmetry, yields the equilibrium value

$$f_1^* = \frac{2}{3}(1-\gamma),$$

which does not depend on p_d. Substituting this value into $\tau_1^*(\cdot)$ and $p_1^*(\cdot)$, we find

$$\tau_1^*(p_d) \triangleq \tau_1^*(f_1^*, f_1^*, p_d) = \frac{5(1-\gamma)}{6} + \gamma p_d$$

and

$$p_1^*(p_d) \triangleq p_1^*(f_1^*, f_1^*, p_d) = \frac{1}{6}\left(4 - 5\gamma + \frac{3}{2+\gamma}\right) + \gamma p_d.$$

Therefore, at the first stage of the game, M commits to the value of p_d which solves

$$\max_{p_d} 2(\tau_1^*(p_d) - f_1^*) D^i(p_1^*(p_d), p_1^*(p_d), p_d) + p_d D^d(p_d, p_1^*(p_d), p_1^*(p_d)),$$

whose first-order condition is

$$\begin{aligned}
&2\frac{\partial \tau_1^*(p_d)}{\partial p_d}\left(\frac{1-\gamma-(1+\gamma)\,p_1^*(p_d)+\gamma\,(p_1^*(p_d)+p_d)}{(1-\gamma)(1+2\gamma)}\right)+\left(\frac{1-\gamma-(1+\gamma)\,p_d+\gamma 2 p_1^*(p_d)}{(1-\gamma)(1+2\gamma)}\right)+\\
&+2(\tau_1^*(p_d)-f_1^*)\left(\frac{\partial p_1^*(p_d)}{\partial p_d}\left(-\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}+\frac{\gamma}{(1-\gamma)(1+2\gamma)}\right)+\frac{\gamma}{(1-\gamma)(1+2\gamma)}\right)+\\
&+p_d\left(-\frac{1+\gamma}{(1-\gamma)(1+2\gamma)}+2\frac{\partial p_1^*(p_d)}{\partial p_d}\frac{\gamma}{(1-\gamma)(1+2\gamma)}\right)=0. \qquad (2.9)
\end{aligned}$$

Solving this equation, we find

$$p_{d,1}^* = p^M,$$

which can be finally substituted into $\tau_1^*(p_d)$ and $p_1^*(p_d)$ to obtain

$$\tau_1^* = f_1^* + \underbrace{\frac{1+2\gamma}{6}}_{(>0)\ M\text{'s mark-up}}$$

and

$$p_1^* = \tau_1^* + \underbrace{\frac{1-\gamma}{6(2+\gamma)}}_{(>0)\ I_i\text{'s mark-up}},$$

with $p_1^* \geq \tau_1^* \geq p^M$.

Equilibrium without platform parity. For any pair of commissions (f_A, f_B) negotiated with the platforms, and any announced value of p_d, M's optimization problem at the final pricing stage is as follows:

$$\max_{\tau_A,\tau_B} \sum_{i=A,B} (\tau_i - f_i) D^i(p_0^*(\tau_i, p_d), p_0^*(\tau_{-i}, p_d), p_d) + p_d D^d(p_d, p_0^*(\tau_A, p_d), p_0^*(\tau_B, p_d)),$$

whose first-order conditions are as in the baseline agency model of Section 2.3.2 and yield

$$\tau_0^*(f_i, p_d) \triangleq \frac{f_i}{2} + \frac{1+\gamma(1-p_0^*+2p_d+\gamma(3p_d-2))}{2(1+\gamma)}, \quad \forall i = A, B.$$

Notice that, differently from the model without commitment, here the access price charged to I_i only depends on the commission negotiated with P_i.

Moving backward at the previous stage of the game, P_i's problem is

$$\max_{f_i} f_i D^i(p_0^*(f_i, p_d), p_0^*(f_0^*, p_d), p_d),$$

where, to ease notation, $p_0^*(f_i, p_d) \triangleq p_0^*(\tau_0^*(f_i, p_d), p_d)$. The first-order condition for this optimization problem, which is analogous to the corresponding one shown in Section 2.3.2, imposing symmetry (i.e., $f_i = f_0^*$), yields the equilibrium value

$$f_0^* = \frac{4(1-\gamma)^2}{8+5\gamma+3\gamma^2},$$

which, also in this case, does not depend on p_d. We can then obtain

$$p_0^*(p_d) \triangleq p_0^*(f_0^*, p_d) = \frac{(1-\gamma)(7+2\gamma)}{8+3\gamma} + \gamma p_d.$$

Therefore, at the first stage of the game, M commits to the value of p_d which solves

$$\max_{p_d} 2(\tau_0^*(f_0^*, p_d) - f_0^*) D^i(p_0^*(p_d), p_0^*(p_d), p_d) + p_d D^d(p_d, p_0^*(p_d), p_0^*(p_d)),$$

whose first-order condition is analogous to equation (2.9), and, also in this case, yields

$$p_{d,0}^* = p^M,$$

which can be finally substituted into $\tau_0^*(f_0^*, p_d)$ and $p_0^*(p_d)$, to obtain

$$\tau_0^* = f_0^* + \underbrace{\frac{4+8\gamma-\gamma^2}{16+6\gamma}}_{M\text{'s mark-up}},$$

and

$$p_0^* = \tau_0^* + \underbrace{\frac{(1-\gamma)^2}{8-5\gamma-3\gamma^2}}_{I_i\text{'s mark-up}},$$

with $p_0^* \geq \tau_0^* \geq p^M$.

Consumer Welfare. Since, in both parity regimes, M commits to set the same price on the direct channel, to evaluate the effects of the parity provision on consumer welfare we only need to compare final prices in the indirect channel. It is easy to see that

$$p_1^* - p_0^* = \frac{(1-\gamma^2)(4+3\gamma)}{6(2+\gamma)(8+3\gamma)} > 0,$$

which concludes the proof. ■

$N \geq 2$ **competing platforms.** We begin by considering I_i's problem in regime $k = 1, 0$:

$$\max_{p_i} D^i\left(p_i, \sum_{j\in\mathcal{I}\setminus\{i\}} p_{j,k}^*\right)(p_i - \tau_i),$$

where, since we are restricting our attention to a symmetric equilibrium,

$$\sum_{j\in\mathcal{I}\setminus\{i\}} p_{j,k}^* = (N-1)p_k^* + p_{d,k}^*.$$

The (standard) first-order condition for the considered problem is

$$\frac{1-\gamma}{1+(N-1)\gamma-N\gamma^2} - \frac{1+(N-1)\gamma}{1+(N-1)\gamma-N\gamma^2}(p_i-\tau_i)+$$

$$+\frac{1-\gamma}{1+(N-1)\gamma-N\gamma^2}\left(\frac{\gamma}{1-\gamma}((N-1)p_k^*+p_{d,k}^*) - \frac{1+(N-1)\gamma}{1-\gamma}p_i\right) = 0. \quad (2.10)$$

Equilibrium with platform parity. When platform parity is in place (i.e., $\tau_i = \tau$ for all $i = 1, \ldots, N$), by solving the system of first-order conditions (2.10), we obtain the equilibrium retail price

$$p_1^*(\tau) \triangleq \left(1 - \frac{1}{2 + (N-1)\gamma}\right)\tau + \frac{1 - \gamma(1 - p_{d,1}^*)}{2 + (N-1)\gamma}.$$

At the final pricing stage, M solves

$$\max_{p_d} \sum_{i=1,\ldots,N} (\tau - f_i) D^i\left(p_1^*(\tau), (N-1)p_1^*(\tau) + p_d\right) + p_d D^d(p_d, Np_1^*(\tau)),$$

whose first-order condition,

$$\frac{\gamma}{1 + (N-1)\gamma - N\gamma^2} \sum_{i=1,\ldots,N} (\tau - f_i) + \frac{1-\gamma}{1 + (N-1)\gamma - N\gamma^2} - \frac{1 + (N-1)\gamma}{1 + (N-1)\gamma - N\gamma^2} p_d +$$

$$+ \frac{\gamma}{1 + (N-1)\gamma - N\gamma^2} N p_1^*(\tau) - \frac{1 + (N-1)\gamma}{1 + (N-1)\gamma - N\gamma^2} p_d = 0,$$

gives a function $p_{d,1}^*(\boldsymbol{f})$, where $\boldsymbol{f} \triangleq (f_1, \ldots, f_N) \in \mathbb{R}^N$. In the previous stage of the game, M sets the (common) access price by solving

$$\max_{\tau} \sum_{i=1,\ldots,N} (\tau - f_i) D^i\left(p_1^*(\tau), (N-1)p_1^*(\tau) + p_{d,1}^*(\boldsymbol{f})\right) + p_{d,1}^*(\boldsymbol{f}) D^d(p_{d,1}^*(\boldsymbol{f}), Np_1^*(\tau)),$$

whose first-order condition, by the Envelope Theorem, is

$$\frac{1-\gamma}{1 + (N-1)\gamma - N\gamma^2}\left(1 - \frac{1 + (N-1)\gamma}{1-\gamma} p_1^*(\tau) + \frac{\gamma}{1-\gamma}((N-1)p_1^*(\tau) + p_{d,1}^*(\boldsymbol{f}))\right) +$$

$$+ \frac{\partial p_1^*(\tau)}{\partial \tau}\left(-\frac{1 + (N-1)\gamma}{1 + (N-1)\gamma - N\gamma^2} + (N-1)\frac{\gamma}{1 + (N-1)\gamma - N\gamma^2}\right)\left(\tau - \frac{1}{N}\sum_{i=1,\ldots,N} f_i\right) +$$

$$+ \frac{\partial p_1^*(\tau)}{\partial \tau} p_{d,1}^*(\boldsymbol{f}) \frac{\gamma}{1 + (N-1)\gamma - N\gamma^2} = 0,$$

and it yields a best-reply function $\tau_1^*(\boldsymbol{f})$.

Finally, we consider P_i's problem. Notice that, when P_i offers a contract f_i to M, it believes that every other platform is offering the equilibrium commission f_1^*. Hence, for every offer f_i, M is expected to set prices $\tau_1^*(f_i, \boldsymbol{f}_{-i,1}^*)$ and $p_{d,1}^*(f_i, \boldsymbol{f}_{-i,1}^*)$, where we denote $\boldsymbol{f}_{-i,1}^* \triangleq (f_1^*, \ldots, f_1^*) \in \mathbb{R}^{N-1}$. Accordingly, intermediaries are expected to charge a retail price $p_1^*(f_i, \boldsymbol{f^*_{-i,1}}) \triangleq p_1^*(\tau_1^*(f_i, \boldsymbol{f^*_{-i,1}}))$. Thus, P_i's problem is as follows:

$$\max_{f_i} f_i D^i \left(p_1^*(f_i, \boldsymbol{f^*_{-i,1}}), (N-1)p_1^*(f_i, \boldsymbol{f^*_{-i,1}}) + p_{d,1}^*(f_i, \boldsymbol{f^*_{-i,1}})\right).$$

By solving the first-order condition,

$$\frac{1-\gamma}{1+(N-1)\gamma - N\gamma^2} - \frac{1+(N-1)\gamma}{1+(N-1)\gamma - N\gamma^2} p_1^*(f_i, \boldsymbol{f^*_{-i,1}}) +$$

$$+\frac{\gamma((N-1)p_1^*(f_i, \boldsymbol{f^*_{-i,1}}) + p_{d,1}^*(f_i, \boldsymbol{f^*_{-i,1}}))}{1+(N-1)\gamma - N\gamma^2} + f_i\left[-\frac{1+(N-1)\gamma}{1+(N-1)\gamma - N\gamma^2}\frac{\partial p_1^*(f_i, \boldsymbol{f^*_{-i,1}})}{\partial f_i}\right] +$$

$$+\frac{\gamma}{1+(N-1)\gamma - N\gamma^2}\left((N-1)\frac{\partial p_1^*(f_i, \boldsymbol{f^*_{-i,1}})}{\partial f_i} + \frac{\partial p_{d,1}^*(f_i, \boldsymbol{f^*_{-i,1}})}{\partial f_i}\right)\Bigg] = 0.$$

and imposing symmetry — i.e., $f_i = f_1^*$ and $p_{d,1}^*(f_1^*, \boldsymbol{f_1^*}) = p_{d,1}^*$ — we obtain the equilibrium values

$$f_1^* = \frac{1}{\Psi_1} N(1-\gamma)(8+\gamma(20(N-1)+\gamma(16-41N+16N^2-4N\gamma^2(N-1)^2+4\gamma(N-1)(1-5N+N^2))))$$

and

$$p_{d,1}^* = p^M + \frac{1}{\Psi_1} N\gamma(1-\gamma)(1+\gamma(N-1)),$$

where

$$\Psi_1 \triangleq 8(1+N) - \gamma(20(1-N^2) - \gamma(16-24N-25N^2+16N^3)+$$
$$+4\gamma^2(1-5N+5N^3-N^4) + 4N(N+1)(N-1)^2\gamma^3).$$

Finally, substituting these values into $\tau_1^*(\cdot)$ and $p_1^*(\cdot)$, we obtain

$$\tau_1^* = f_1^* + \frac{1}{2\Psi_1}(8-\gamma(20-28N+\gamma(-16+58N-36N^2)+\gamma^2(4-36N+63N^2-20N^3))+$$
$$-2\gamma^3 N(N-1)(3-12N+2N^2) + 4\gamma^4 N^2(N-1)^2)$$

and

$$p_1^* = \tau_1^* + \frac{2}{\Psi_1}(1-\gamma)^2(1+\gamma(N-1))(1+N\gamma).$$

Equilibrium without platform parity. In the absence of platform parity, the first-order condition (2.10) yields I_i's best-reply function

$$p_0^*(\tau_i) \triangleq \frac{\tau_i}{2} + \frac{1+\gamma((N-1)p_0^* + p_{d,0}^* - 1)}{2(1+\gamma(N-1))}.$$

Hence, at the final pricing stage, M's optimization problem is

$$\max_{p_d} \sum_{i=1,\ldots,N} (\tau_i - f_i) D^i \left(p_0^*(\tau_i), (N-1)p_0^*(\tau_i) + p_d\right) + p_d D^d(p_d, N p_0^*(\tau_i)),$$

whose first-order condition is

$$\frac{\gamma}{1+(N-1)\gamma - N\gamma^2} \sum_{i=1,\ldots,N} (\tau_i - f_i) - \frac{1+(N-1)\gamma}{1+(N-1)\gamma - N\gamma^2} p_d +$$

$$+\frac{1-\gamma}{1+(N-1)\gamma - N\gamma^2} \left(1 - \frac{1+(N-1)\gamma}{1-\gamma} p_d + \frac{\gamma}{1-\gamma} N p_0^*(\tau_i)\right) = 0.$$

Letting $p_{d,0}^*(\boldsymbol{f})$ denote the solution to the above equation, M's problem in the previous stage of the game is as follows:

$$\max_{\tau_1,\ldots,\tau_N} \sum_{i=1,\ldots,N} (\tau_i - f_i) D^i \left(p_0^*(\tau_i), \sum_{j \in \mathcal{I} \setminus \{i,d\}} p_0^*(\tau_j) + p_{d,0}^*(\boldsymbol{f})\right) + p_{d,0}^*(\boldsymbol{f}) D^d \left(p_{d,0}^*(\boldsymbol{f}), \sum_{i \in \mathcal{I} \setminus \{d\}} p_0^*(\tau_i)\right).$$

By the Envelope Theorem, the first-order condition with respect to τ_i is

$$\frac{1-\gamma}{1+(N-1)\gamma - N\gamma^2} \left(1 - \frac{1+(N-1)\gamma}{1-\gamma} p_0^*(\tau_i) + \frac{\gamma}{1-\gamma} \left(\sum_{j \in \mathcal{I} \setminus \{i,d\}} p_0^*(\tau_j) + p_{d,0}^*(\boldsymbol{f})\right)\right) +$$

$$-\frac{1+(N-1)\gamma}{1+(N-1)\gamma - N\gamma^2} (\tau_i - f_i) \frac{\partial p_0^*(\tau_i)}{\partial \tau_i} + \frac{\gamma}{1+(N-1)\gamma - N\gamma^2} \frac{\partial p_0^*(\tau_i)}{\partial \tau_i} \sum_{j \in \mathcal{I} \setminus \{i,d\}} (\tau_j - f_j) +$$

$$+\frac{\gamma}{1+(N-1)\gamma - N\gamma^2} p_{d,0}^*(\boldsymbol{f}) \frac{\partial p_0^*(\tau_i)}{\partial \tau_i} = 0.$$

By solving the system of these first-order conditions, we find M's best-reply functions $\tau_{i,0}^*(\boldsymbol{f})$.

Finally, we analyze platforms' problem. When P_i offers a contract f_i to M, it believes that every other platform is offering the equilibrium commission f_0^*. Hence, for every offer f_i, denoting $\boldsymbol{f^*_{-i,0}} \triangleq (f_0^*, \ldots, f_0^*) \in \mathbb{R}^{N-1}$, M is expected to set: a price $p_{d,0}^*(f_i, \boldsymbol{f^*_{-i,0}})$ on the direct sale channel; an access price $\tau_{i,0}^*(f_i, \boldsymbol{f^*_{-i,0}})$ on P_i; an access price $\tau_{-i,0}^*(f_i, \boldsymbol{f^*_{-i,0}})$ on every other platform. Hence, P_i's problem is as follows:

$$\max_{f_i} f_i D^i \left(p_{i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}), (N-1)p_{-i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}) + p_{d,0}^*(f_i, \boldsymbol{f^*_{-i,0}})\right)$$

where $p_{i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}) \triangleq p_0^*(\tau_{i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}))$ and $p_{-i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}) \triangleq p_0^*(\tau_{-i,0}^*(f_i, \boldsymbol{f^*_{-i,0}}))$.

Solving the first-order condition

$$\frac{1-\gamma}{1+(N-1)\gamma-N\gamma^2}-\frac{1+(N-1)\gamma}{1+(N-1)\gamma-N\gamma^2}p^*_{i,0}(f_i,\boldsymbol{f^*_{-i,0}})+$$

$$+\frac{\gamma((N-1)p^*_{-i,0}(f_i,\boldsymbol{f^*_{-i,0}})+p^*_{d,0}(f_i,\boldsymbol{f^*_{-i,0}}))}{1+(N-1)\gamma-N\gamma^2}+f_i\Bigg[-\frac{1+(N-1)\gamma}{1+(N-1)\gamma-N\gamma^2}\frac{\partial p^*_{i,0}(f_i,\boldsymbol{f^*_{-i,0}})}{\partial f_i}+$$

$$+\frac{\gamma}{1+(N-1)\gamma-N\gamma^2}\left((N-1)\frac{\partial p^*_{-i,0}(f_i,\boldsymbol{f^*_{-i,0}})}{\partial f_i}+\frac{\partial p^*_{d,0}(f_i,\boldsymbol{f^*_{-i,0}})}{\partial f_i}\right)\Bigg]=0,$$

and imposing symmetry (i.e., $f_i=f^*_0$, $\tau^*_{i,0}(f^*,\boldsymbol{f^*_{-i,0}})=\tau^*_0$, $p^*_{i,0}(f^*,\boldsymbol{f^*_{-i,0}})=p^*_0$, and $p^*_{d,0}(f^*,\boldsymbol{f^*_{-i,0}})=p^*_{d,0}$) we obtain the equilibrium values

$$f^*_0=\frac{1}{\Psi_0}(2(1-\gamma)(1+\gamma(N-1))(2+\gamma(3N-2))(8-\gamma(8-(8-9\gamma)N)),$$

$$p^*_{d,0}=p^M+\frac{1}{2\Psi_0}N\gamma(1+\gamma(N-1))(8-\gamma(24-20N+\gamma(-25+41N-12N^2+3\gamma(N-1)(5N-4))),$$

$$\tau^*_0=f^*_0+\frac{1}{2\Psi_0}(32-\gamma(96-144N-4\gamma(27-89N+60N^2)+4\gamma^2(19-82N+121N^2-44N^3)+$$
$$+\gamma^3(-41+161N-334N^2+284N^3-48N^4)))+3\gamma^4(N-1)(-4+9N-18N^2+20N^3),$$

$$p^*_0=\tau^*_0+\frac{1}{\Psi_0}(1-\gamma)(1+N\gamma)(8-\gamma(24-20N+\gamma(-25+41N-12N^2+3\gamma(N-1)(5N-4)))),$$

where

$$\Psi_0\triangleq 64+\gamma(-216+248N-9N\gamma^4(N-1)^2(5N-4)+4\gamma(67-169N+85N^2)+$$
$$-\gamma^2(155-662N+733N^2-192N^3)-6\gamma^3(N-1)(6-40N+47N^2-6N^3)).$$

M's incentive to delist. As in the model with two platforms, assuming that M is allowed to delist from platforms offering excessively high commissions, we must check whether, in each regime k, f^*_k satisfies M's participation constraint — i.e., whether M, who has been offered the commission f^*_k from a platform P_i, does not find it profitable to decline that offer (hence, to shut down that channel), while still accepting the same offer made by the remaining $N-1$ platforms.[21]

To this end, we should determine M's *deviation profit*, that is the profit that M gets by serving $N-1$ platforms, and compare it with the *equilibrium profit*, obtained by accepting all platforms' offers.

[21]Notice that, since we are considering a contract equilibrium, we disregard the possibility of multilateral deviations by M — i.e., that M can refuse more than one offer.

Let us examine the regime with platform parity. Considering M's best-reply functions $\tau_1^*(f_i, \boldsymbol{f^*_{-i,1}})$ and $p^*_{d,1}(f_i, \boldsymbol{f^*_{-i,1}})$, computed assuming that the number of active platforms is $N-1$, and substituting the equilibrium value (with N platforms) f_1^*, we obtain that M's optimal choices after refusing P_i's offer are

$$\tilde{p}_{d,1} = p^M + \frac{\gamma(1-\gamma)(2+\gamma(N-1))(N-1)(1+N\gamma)}{(2+\gamma(N-2))\Psi_1},$$

$$\tilde{\tau}_1 = f_1^* + \frac{1}{2(2+\gamma(N-2))\Psi_1}(16-\gamma(56-64N-4\gamma(17-48N+25N^2)+\gamma^2(34-204N+256N^2-76N^3)+$$
$$-\gamma^3(6-90N+224N^2-159N^3+28N^4)+2\gamma^4N(N-1)(7-31N+20N^2-2N^3)+4N^2(N-2)(N-1)^2\gamma^5)).$$

Accordingly, intermediaries' optimal price becomes[22]

$$\tilde{p}_1 = \left(1 - \frac{1}{2+(N-2)\gamma}\right)\tau + \frac{1-\gamma(1-\tilde{p}_{d,1})}{2+(N-2)\gamma}.$$

From these values, it is easy to determine M's *deviation profit*. Finally, in Figure 2.15, we show that M's participation constraint is satisfied when all platforms offer f_1^*.[23] It can be easily checked that, a fortiori, M does not find it profitable to decline P_i's offer f_0^* in the game without platform parity. ■

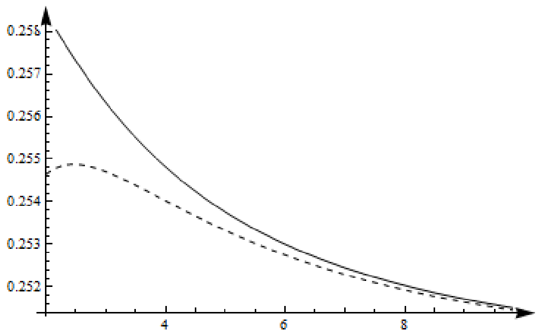

Fig. 2.15: M's equilibrium profit (continuous line) and deviation profit (dashed line) in the game with platform parity agreement as a function of $N \geq 2$. Parameter's value: $\gamma = .5$.

Proof of Proposition 2.11. Following the same steps of the analysis of Sections 2.3.1 and 2.3.2, we can characterize the equilibrium of the game under the two parity regimes, for any given value of $c \in [0, 1-\gamma]$.

[22]Recall that we assumed that intermediaries observe M's listing decision.
[23]By using other numerical examples, we checked that M's *deviation profit* is lower than its *equilibrium profit* for many different values of γ.

Specifically, if a platform parity agreement is in place, platform's commissions are

$$f_1^*(c) = \frac{(4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4)(1-\gamma-c)}{6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4};$$

M sets price

$$p_{d,1}^*(c) = p^M + \frac{\gamma(1+\gamma)(1-\gamma-c)}{2(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4)}$$

on the direct sale channel, and charges the access price

$$\tau_1^*(c) = f_1^*(c) + \frac{2+9\gamma+11\gamma^2-6\gamma^3-13\gamma^4-4\gamma^5-c(2+5\gamma+\gamma^2-4\gamma^3-2\gamma^4)}{2+9\gamma+11\gamma^2-6\gamma^3-13\gamma^4-4\gamma^5}$$

to the intermediaries, which in turn set

$$p_1^*(c) = \tau_1^*(c) + \frac{(1-\gamma)(1+\gamma)(1+2\gamma)(1-\gamma-c)}{2\left(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4\right)}.$$

In the absence of platform parity, platforms' commissions are

$$f_0^*(c) = \frac{8(1+\gamma)(4+4\gamma-9\gamma^2)(1-c-\gamma)}{64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4};$$

M's price on the direct channel is

$$p_{d,0}^*(c) = p^M + \frac{\gamma(1+\gamma)(8-9\gamma^2)(1-c-\gamma)}{(1-\gamma)(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)},$$

and the access price is

$$\tau_0^*(c) = f_0^*(c) + \frac{32+96\gamma-28\gamma^2-248\gamma^3-5\gamma^4+153\gamma^5-2c(1+\gamma)(8-9\gamma^2)(2+2\gamma-3\gamma^2)}{2(1-\gamma)(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)}.$$

Finally, the intermediaries' final price is

$$p_0^*(c) = \tau_0^*(c) + \frac{(1+2\gamma)(8-9\gamma^2)(1-c-\gamma)}{(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)}.$$

Clearly, in both parity regimes, equilibrium commissions and prices are all decreasing in c.

M's profits in the two regimes are as follows

$$\pi_1^M(c) \triangleq \pi_1^M - \frac{c(2-2\gamma-c)(1+\gamma)^2(4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4)}{4(6+15\gamma-\gamma^2-15\gamma^3-6\gamma^4)^2},$$

$$\pi_0^M(c) \triangleq \pi_0^M - \frac{c(2-2\gamma-c)(1+\gamma)^2(8-9\gamma^2)^2(4+4\gamma-9\gamma^2)}{(1-\gamma)^2(64+152\gamma-28\gamma^2-171\gamma^3-54\gamma^4)^2},$$

where, for every $k = 1, 0$, π_k^M is M's profit when $c = 0$. It is straightforward to see that, in both parity regimes, M's profit is decreasing in c and it can be easily checked that, for every $c \in [0, 1-\gamma]$, $\pi_1^M(c) < \pi_0^M(c)$.
Platforms' equilibrium profits are

$$\pi_{i,1}^P(c) \triangleq \frac{(1-c-\gamma)^2}{(1-\gamma)^2}\pi_{i,1}^P,$$

$$\pi_{i,0}^P(c) \triangleq \frac{(1-c-\gamma)^2}{(1-\gamma)^2}\pi_{i,0}^P,$$

where, for every $k = 1, 0$, $\pi_{i,k}^P$ is a platform's profit when $c = 0$. It is straightforward to see that, in both parity regimes, platforms' profits are decreasing in c and that $\pi_{i,1}^P(c) > \pi_{i,0}^P(c)$ if and only if $\pi_{i,1}^P > \pi_{i,0}^P$. ■

Proof of Proposition 2.12. For all $\gamma > \hat{\gamma}$: (i) $\pi_0^P(0) < \pi_1^P(0)$, (ii) $\pi_1^P(c)$ is decreasing in c, and (iii) $\pi_0^P(0) > \pi_1^P(1-\gamma) = 0$. This shows the existence of a threshold $\bar{c} \in (0, 1-\gamma)$ such that $\pi_0^P(0) > \pi_1^P(c)$, for all $c > \bar{c}$. ■

Appendix 2.B. Second-order conditions

In this Appendix, we derive the second-order conditions for the maximization problems stated in the main analysis.

Multi-product monopolist. The Hessian matrix of the optimization problem of the multi-product monopolist is

$$H(p_A, p_B, p_d) = \begin{pmatrix} -\frac{2(1+\gamma)}{(1-\gamma)(1+2\gamma)} & \frac{2\gamma}{(1-\gamma)(1+2\gamma)} & \frac{2\gamma}{(1-\gamma)(1+2\gamma)} \\ \frac{2\gamma}{(1-\gamma)(1+2\gamma)} & -\frac{2(1+\gamma)}{(1-\gamma)(1+2\gamma)} & \frac{2\gamma}{(1-\gamma)(1+2\gamma)} \\ \frac{2\gamma}{(1-\gamma)(1+2\gamma)} & \frac{2\gamma}{(1-\gamma)(1+2\gamma)} & -\frac{2(1+\gamma)}{(1-\gamma)(1+2\gamma)} \end{pmatrix}$$

whose eigenvalues are

$$\lambda_1 = \lambda_2 = -\frac{2}{1-\gamma} < 0,$$

$$\lambda_3 = -\frac{2}{1+2\gamma} < 0.$$

Hence, second-order conditions are satisfied for all $\gamma \in (0, 1)$.

The wholesale benchmark. It can be easily seen that, in both parity regimes, second-order conditions of intermediaries' and M's problems at the final pricing stage, as well as the second-order condition of platforms' problem, are satisfied for all $\gamma \in (0, 1)$. Second-order conditions of M's maximization problems concerning wholesale prices are as follows.

- In the regime with platform parity, we have

$$\frac{\partial^2 \pi^M}{\partial t^2} = -\frac{2(2+\gamma)(4+7\gamma-10\gamma^2-4\gamma^3)}{(1-\gamma^2)(4+\gamma)^2(1+2\gamma)}.$$

  Therefore, the second-order condition is satisfied if and only if

$$\gamma < \overline{\gamma} \approx 0.86.$$

- In the regime without platform parity, we have

$$\frac{\partial^2 \pi^M}{\partial t_i^2} = -\frac{(4-\gamma)(4+9\gamma)}{32(1-\gamma^2)(1+2\gamma)} < 0,$$

  hence, the second-order condition is satisfied for all $\gamma \in (0, 1)$.

The agency model. In both parity regimes, the second-order conditions of M's and the intermediaries' problems at the final pricing stage are satisfied for all $\gamma \in (0, 1)$. The second-order conditions of M's maximization problems concerning access prices and of platforms' problems are as follows.

- In the regime with platform parity, we have

$$\frac{\partial^2 \pi^M}{\partial \tau^2} = -\frac{2(4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4)}{(2+\gamma)^2(1+2\gamma-\gamma^2-2\gamma^3)}$$

  which is negative for every

$$\gamma < \frac{1}{8}\left(\sqrt{5}+\sqrt{62+22\sqrt{5}}-5\right) \approx 0.93.$$

  Finally, the second-order condition for the platforms' problem is satisfied if and only if the above condition holds true, since

$$\frac{\partial^2 \pi_i^P}{\partial f_i^2} = -\frac{(1+\gamma)^2}{(4+10\gamma-\gamma^2-10\gamma^3-4\gamma^4)}.$$

- In the regime without platform parity, we have to consider the following Hessian matrix:

$$H(\tau_A, \tau_B) = \begin{pmatrix} -\frac{8+(16-\gamma)\gamma}{8(1+2\gamma)(1-\gamma^2)} & \frac{\gamma(8+17\gamma)}{8(1+2\gamma-\gamma^2-2\gamma^3)} \\ \frac{\gamma(8+17\gamma)}{8(1+2\gamma-\gamma^2-2\gamma^3)} & -\frac{8+(16-\gamma)\gamma}{8(1+2\gamma)(1-\gamma^2)} \end{pmatrix}$$

whose eigenvalues are

$$\lambda_1 = -\frac{1}{1+\gamma} < 0, \quad \lambda_2 = \frac{1}{8}\left(\frac{1}{1-2\gamma^2+\gamma} - \frac{9}{1+\gamma}\right).$$

Thus, the second-order condition is satisfied if and only if

$$\lambda_2 < 0 \quad \iff \quad \gamma < \frac{2(1+\sqrt{10})}{9} \approx 0.92.$$

Finally, under this condition, also the second-order condition of the platforms' problem is satisfied, since

$$\frac{\partial^2 \pi_i^P}{\partial f_i^2} = \frac{1}{9\gamma^2 - 4(1+\gamma)} - \frac{1}{4(1-\gamma)} < 0 \quad \forall \gamma < 0.92.$$

Appendix 2.C. Additional results

In this Appendix, we first show that our main results, both in the wholesale and the agency model, also apply when considering asymmetric product differentiation between the direct and the indirect channel, then analyze in detail some of the extensions discussed in the paper. Proofs are omitted for brevity and available upon request.

Asymmetric product differentiation. We now show that our main results, both in the wholesale and the agency model, are valid if we assume that the products distributed in the indirect channel are perceived by the consumers as more close substitutes relative to the product distributed in the direct channel.

To this aim, consider a representative consumer whose utility function is

$$U(\cdot) \triangleq \sum_{j=A,B,d} q_j - \frac{1}{2} \sum_{j=A,B,d} q_j^2 - \gamma q_A q_B - \delta(q_A + q_B) q_d - \sum_{j=A,B,d} p_j q_j + m,$$

where m is the utility from income.

Standard techniques then yield the (direct) demand functions

$$q_i = D^i(p_i, p_{-i}, p_d) = \frac{(1-\delta)(1-\gamma) - (1-\delta^2)p_i + (\gamma-\delta^2)p_{-i} + \delta(1-\gamma)p_d}{(1-2\delta^2+\gamma)(1-\gamma)}, \quad \forall i = A, B.$$

and

$$q_d = D^d(p_d, p_A, p_B) = \frac{1 - 2\delta + \gamma + \delta(p_A + p_B) - (1+\gamma)p_d}{1 - 2\delta^2 + \gamma}.$$

We assume that $0 < \delta \leq \gamma < 1$ in order to guarantee that every intermediary faces stronger competition from its rival than the monopolist (via the direct channel).

Clearly, the equilibrium analysis is as in the baseline model. First, it is easy to find that the multi-product monopolist still finds it optimal to set $p^M = \frac{1}{2}$ on all channels.[24]

In what follows, we consider first the wholesale model and then the agency model.

The wholesale model. With platform parity:[25]

(i) the monopolist sets

$$p^*_{d,1} = p^M + \underbrace{\frac{3(1-\delta)(1-\gamma)\delta}{2(1-2\delta^2+\gamma)(4-\delta^2-3\gamma)}}_{(>0)\text{ Channel externality}}$$

and charges

$$t^*_1 = p^M + \frac{3(1-\delta)(1-\gamma)\delta^2}{2(1-2\delta^2+\gamma)(4-\delta^2-3\gamma)}$$

to the platforms.

(ii) The platforms charge

$$w^*_1 = t^*_1 + \underbrace{\frac{(1-\gamma)(1-\delta)}{4-3\gamma-\delta^2}}_{(>0)\ P_i\text{'s mark-up}}$$

to the intermediaries.

(iii) The intermediaries set

$$p^*_1 = w^*_1 + \underbrace{\frac{(1-\gamma)(1-\delta)}{2(4-3\gamma-\delta^2)}}_{(>0)\ I_i\text{'s mark-up}}.$$

Without platform parity:[26]

[24]Details are available upon request.

[25]It can be easily checked that all second-order conditions of players' optimization problems are satisfied if and only if

$$\gamma > \frac{1 + 13\delta^2 - 5\delta^4 - \sqrt{(1-\delta^2)^3(49-13\delta^2)}}{3(2+\delta^2)}.$$

[26]It can be easily checked that all second-order conditions of players' optimization problems are satisfied for all $\delta \leq \gamma$.

(i) the monopolist sets

$$p_{d,0}^{*} = p^{M} + \underbrace{\frac{3\delta\left(1+\delta\right)\left(1-\delta\right)^{2}}{(8-3\gamma-5\delta^{2})(1+\gamma-2\delta^{2})}}_{(>0)\ \text{Channel externality}}$$

and charges

$$t_{0}^{*} = p^{M} + \frac{3\left(1-\delta\right)\left(\gamma\left(1+\gamma\right)+\delta^{2}\left(1-3\gamma\right)\right)}{2(\gamma-2\delta^{2}+1)(8-3\gamma-5\delta^{2})}$$

to the platforms.

(ii) The platforms charge

$$w_{0}^{*} = t_{0}^{*} + \underbrace{\frac{2\left(1-\delta\right)\left(1-\gamma\right)}{8-3\gamma-5\delta^{2}}}_{(>0)\ P_i\text{'s mark-up}}$$

to the intermediaries.

(iii) The intermediaries set

$$p_{0}^{*} = w_{0}^{*} + \underbrace{\frac{\left(1-\gamma\right)\left(1-\delta\right)}{8-3\gamma-5\delta^{2}}}_{(>0)\ I_i\text{'s mark-up}}.$$

The comparison between the equilibrium prices in the two regimes is as follows:

$$p_{1}^{*} - p_{0}^{*} = -\frac{3(1-\delta)(\gamma-\delta^{2})(1+\gamma+2\delta^{2}(1-2\gamma))}{2(1-2\delta^{2}+\gamma)(4-\delta^{2}-3\gamma)(8-5\delta^{2}-3\gamma)} < 0,$$

$$p_{d,1}^{*} - p_{d,0}^{*} = -\frac{3\delta(1-\delta)(\gamma-\delta^{2})(5-2\delta^{2}-3\gamma)}{2(1-2\delta^{2}+\gamma)(4-\delta^{2}-3\gamma)(8-5\delta^{2}-3\gamma)} < 0.$$

Hence, for every $\delta \leq \gamma$, platform parity benefits consumers in the wholesale model.

As for the equilibrium profits, in the presence of platform parity,

$$\pi_{i,1}^{I} = \frac{(1-\delta)^{3}(1+\delta)(1-\gamma)}{4\left(1-2\delta^{2}+\gamma\right)\left(4-\delta^{2}-3\gamma\right)^{2}},$$

$$\pi_{i,1}^{P} = \frac{(1-\delta)^{3}(1+\delta)(1-\gamma)}{2\left(1-2\delta^{2}+\gamma\right)\left(4-\delta^{2}-3\gamma\right)^{2}}$$

and

$$\pi_{1}^{M} = \frac{1}{4} + \frac{(1-\delta)^{2}\left(2(1+\gamma)(4-3\gamma)-\delta^{2}(7-\gamma)(5-3\gamma)+2\delta^{4}(11-5\gamma)-4\delta^{6}\right)}{4\left(1-2\delta^{2}+\gamma\right)^{2}\left(4-\delta^{2}-3\gamma\right)^{2}}.$$

In the game without platform parity, we have

$$\pi_{i,0}^{I}=\frac{(1-\delta)^3(1+\delta)(1-\gamma)}{\left(1-2\delta^2+\gamma\right)\left(8-5\delta^2-3\gamma\right)^2},$$

$$\pi_{i,0}^{P}=\frac{2(1-\delta)^3(1+\delta)(1-\gamma)}{\left(1-2\delta^2+\gamma\right)\left(8-5\delta^2-3\gamma\right)^2}$$

and

$$\pi_{0}^{M}=\frac{\Gamma_0}{4\left(1-2\delta^2+\gamma\right)^2\left(8-5\delta^2-3\gamma\right)^2},$$

where

$$\Gamma_0 \triangleq 200\delta^7+4\delta^6(13\gamma-72)-16\delta^5(4\gamma+29)+\delta^4((26-59\gamma)\gamma+669)+8\delta^3(16\gamma+41)+ \\ -2\delta^2(3\gamma^3-68\gamma^2+113\gamma+234)-64\delta(\gamma+1)+(\gamma+1)\left(9\gamma^3-39\gamma^2+16\gamma+96\right).$$

Comparing these expression yields

$$\pi_{i,1}^{I}-\pi_{i,0}^{I}=\frac{3(1-\delta)^3(1+\delta)(1-\gamma)\left(\gamma-\delta^2\right)\left(16-7\delta^2-9\gamma\right)}{4\left(1-2\delta^2+\gamma\right)\left(4-\delta^2-3\gamma\right)^2\left(8-5\delta^2-3\gamma\right)^2}>0,$$

$$\pi_{i,1}^{P}-\pi_{i,0}^{P}=\frac{3(1-\delta)^3(1+\delta)(1-\gamma)\left(\gamma-\delta^2\right)\left(16-7\delta^2-9\gamma\right)}{2\left(1-2\delta^2+\gamma\right)\left(4-\delta^2-3\gamma\right)^2\left(8-5\delta^2-3\gamma\right)^2}>0,$$

$$\pi_{1}^{M}-\pi_{0}^{M}=\frac{\Gamma_1}{4\left(1-2\delta^2+\gamma\right)^2\left(4-\delta^2-3\gamma\right)^2\left(8-5\delta^2-3\gamma\right)^2},$$

where

$$\Gamma_1 \triangleq 9(1-\delta)^2\left(\gamma-\delta^2\right)\left(2\delta^6(9-11\gamma)-\delta^4(55\gamma^2-148\gamma+81)-\delta^2(3\gamma^3-80\gamma^2+161\gamma-72)+2\gamma(1+\gamma)(4-3\gamma)\right),$$

and it can be shown that $\Gamma_1>0$ for $\delta\leq\gamma$.

Therefore, also in the presence of asymmetric product differentiation, platform parity increases profits of all firms in the considered industry.

The agency model. In the presence of a platform parity provision:[27]

[27] It can be easily checked that all second-order conditions of players' optimization problems are satisfied for $\gamma>\frac{1+\delta^2-\delta^4-\sqrt{(1-\delta)^3(9-5\delta^2)}}{2-\delta^2}$. Moreover, it can be graphically shown that M's participation constraint is not binding in equilibrium. Details are available upon request.

(i) each platform charges

$$f_1^* = \underbrace{\frac{(1-\delta)(4\delta^6 - 2\delta^4(7-\gamma) - 2(2-\gamma)(1+\gamma) + \delta^2(3-\gamma)(5+\gamma))}{6\delta^6 - 3\delta^4(7-\gamma) - 3(2-\gamma)(1+\gamma) + 2\delta^2(11-\gamma-\gamma^2))}}_{(>0)\ P_i\text{'s mark-up}}$$

to the monopolist.

(ii) The monopolist sets

$$p_{d,1}^* = p^M + \underbrace{\frac{(1-\delta)^2\delta(1+\delta)(1-\gamma)}{2(3\delta^4(7-\gamma) + 3(2-\gamma)(1+\gamma) - 2\delta^2(11-\gamma-\gamma^2) - 6\delta^6)}}_{\text{Channels externality}}$$

and charges

$$\tau_1^* = f_1^* + \underbrace{\frac{1}{2} - \frac{(1-\delta)\left(\delta^2+\gamma-2\right)\left(4\delta^4 - \delta^2(\gamma+7) + 2(\gamma+1)\right)}{2\left(6\delta^6 + 3\delta^4(\gamma-7) - 2\delta^2\left(\gamma^2+\gamma-11\right) + 3(\gamma-2)(\gamma+1)\right)}}_{(>0)\ M\text{'s mark-up}}$$

to the intermediaries.

(iii) The intermediaries set

$$p_1^* = \tau_1^* + \underbrace{\frac{(\delta-1)^2(\delta+1)(1-\gamma)\left(2\delta^2-\gamma-1\right)}{2\left(6\delta^6 + 3\delta^4(\gamma-7) - 2\delta^2\left(\gamma^2+\gamma-11\right) + 3(\gamma-2)(\gamma+1)\right)}}_{(>0)\ I_i\text{'s mark-up}}.$$

In the game without platform parity:[28]

(i) each platform charges

$$f_0^* = \underbrace{\frac{8(1-\delta)^2(1+\delta)(1-\gamma)(4(1+\gamma) - 9\delta^2)}{8(1+\gamma)(8-5\gamma) + 3\delta^4(73-21\gamma) - \delta^2(228-\gamma(37+41\gamma)) - 54\delta^6}}_{(>0)\ P_i\text{'s mark-up}}$$

to the monopolist.

[28] It can be easily checked that all second-order conditions of players' optimization problems are satisfied for $\gamma > \frac{9}{4}\delta^2 - 1$. Moreover, it can be graphically shown that M's participation constraint is not binding in equilibrium. Details are available upon request.

(ii) The monopolist sets

$$p_{d,0}^* = p^M + \underbrace{\frac{(1-\delta)^2\delta(1+\delta)(8-9\delta^2)}{8(1+\gamma)(8-5\gamma)+3\delta^4(73-21\gamma)-\delta^2(228-\gamma(37+41\gamma))-54\delta^6}}_{\text{Channels externality}}$$

and charges

$$\tau_0^* = f_0^* + \underbrace{\frac{1}{2}\left[1-\frac{(1-\delta)\left(9\delta^4(11\gamma-9)+\delta^2(112-\gamma(41\gamma+105))+8\left(5\gamma^2+\gamma-4\right)\right)}{2\left(54\delta^6+3\delta^4(21\gamma-73)+\delta^2(228-\gamma(41\gamma+37))+8(\gamma+1)(5\gamma-8)\right)}\right]}_{(>0)\ M_i\text{'s mark-up}}$$

to the intermediaries.

(iii) The intermediaries set

$$p_0^* = \tau_0^* + \underbrace{\frac{(1-\delta)\left(8-9\delta^2\right)(1-\gamma)\left(2\delta^2-\gamma-1\right)}{54\delta^6+3\delta^4(21\gamma-73)+\delta^2(228-\gamma(41\gamma+37))+8(\gamma+1)(5\gamma-8)}}_{(>0)\ I_i\text{'s mark-up}}.$$

Figure 2.16 shows that the results provided in Section 2.3.3, concerning the comparison of equilibrium commissions and prices in the two parity regimes, are robust when asymmetric (horizontal) product differentiation is taken into account. Specifically, for every $\delta \leq \gamma$, platforms' commissions are higher in the presence of platform parity, whereas the price on the direct channel is lower under platform parity. Access prices and, in turn, final prices set by intermediaries, are higher when platform parity is in place, unless both intra-channel and inter-channel competition are sufficiently fierce — i.e., both γ and δ are sufficiently high. As a consequence, in order for the provision to be beneficial to consumers, for every relatively high value of γ, also inter-channel competition (as measured by δ) must be sufficiently intense (see Figure 2.17, panel (d)).

As for firms' equilibrium profits, Figure 2.17 (panels (a),(b),(c)) shows that, in line with the results proved in the paper for the case in which $\delta = \gamma$, if products are asymmetrically differentiated, the introduction of a platform parity agreement unambiguously lowers M's and the intermediaries' profits. Interestingly, platforms are better off when the parity provision is in place for every value of δ, provided that γ is relatively high.

Two-part tariffs in the wholesale model. We now extend the wholesale model by assuming that M offers to each platform P_i a two-part tariff (t_i, T_i) — i.e., a contract specifying a unit wholesale price t_i and a fixed fee T_i. Thus, the total payment that P_i makes to M is given by $t_i D^i(p_i, p_{-i}, p_d) + T_i$. However, as in Calzolari et al. (2018), we

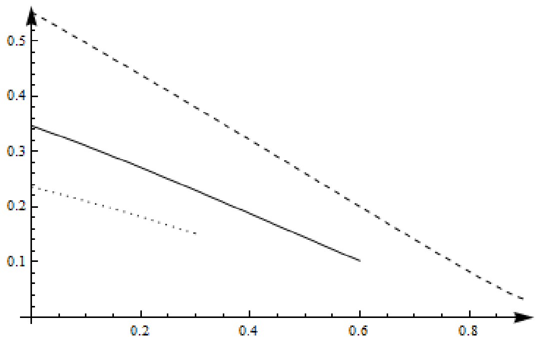

(a) *Difference* $f_1^* - f_0^*$ *as a function of* δ.

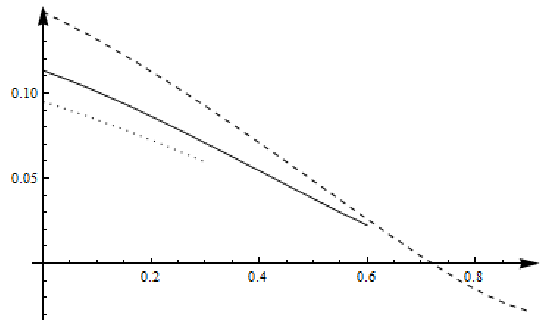

(b) *Difference* $\tau_1^* - \tau_0^*$ *as a function of* δ.

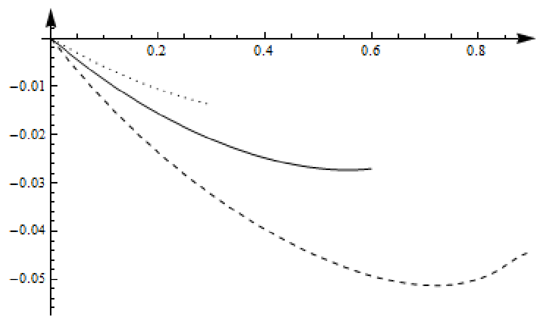

(c) *Difference* $p_{d,1}^* - p_{d,0}^*$ *as a function of* δ.

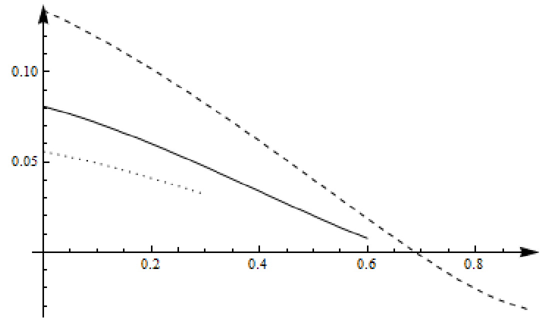

(d) *Difference* $p_1^* - p_0^*$ *as a function of* δ.

Fig. 2.16: Equilibrium values as functions of δ. Parameter's values: $\gamma = 0.3$ (dotted lines), $\gamma = 0.6$ (continuous lines), $\gamma = 0.9$ (dashed lines).

assume that it is costly to extract the platforms' rents by means of fixed fees. Specifically, by charging a fixed fee $T_i > 0$, M gains T_i but P_i loses $(1+\mu)T_i$, with $\mu > 0$.

By the same token, each platform P_i in turn offers a two-part tariff (w_i, W_i) to I_i — i.e., the total payment that I_i makes to P_i is given by $w_i D^i(p_i, p_{-i}, p_d) + W_i$. Again, fixed fees are costly, in the sense that, by charging a fixed fee $W_i > 0$, P_i gains W_i but I_i loses $(1+\mu)W_i$.

We are able to prove the following result.

Proposition 2.C.1. *If all contracts are specified as two-part tariffs, then, in the presence of a platform parity agreement, equilibrium prices are*

$$p_{d,1}^* = p^M + \frac{\gamma(1 - 2\mu + 3\mu^2)}{2(1+2\gamma)(-1-\gamma+(4+\gamma)\mu^2)}$$

and

$$p_1^* = p^M + \frac{(1-\gamma)(3\mu^2 + \gamma(1+\mu)(-1+3\mu))}{2(1+2\gamma)(-1-\gamma+(4+\gamma)\mu^2)}$$

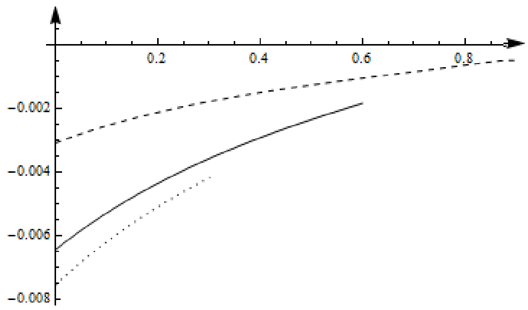

(a) *Difference* $\pi^I_{i,1} - \pi^I_{i,0}$ *as a function of* δ.

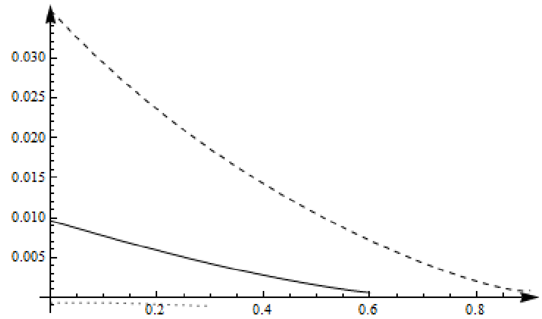

(b) *Difference* $\pi^P_{i,1} - \pi^P_{i,0}$ *as a function of* δ.

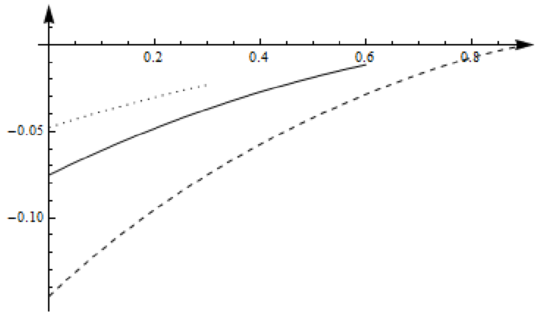

(c) *Difference* $\pi^M_1 - \pi^M_0$ *as a function of* δ.

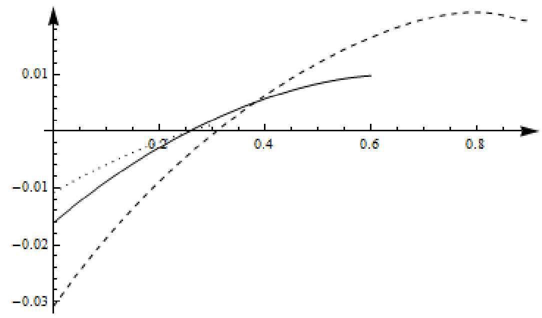

(d) *Difference between the consumer surplus in the presence and in the absence of a platform parity agreement as a function of* δ.

Fig. 2.17: Comparison of equilibrium profits and consumer surplus in the agency model. Parameter's values: $\gamma = 0.3$ (dotted lines), $\gamma = 0.6$ (continuous lines), $\gamma = 0.9$ (dashed lines).

whereas, without platform parity,

$$p^*_{d,0} = p^M + \frac{\gamma(1+\gamma)(1-2\mu+3\mu^2)}{(1+2\gamma)(-2+8\mu^2+\gamma(-1-2\mu+5\mu^2))}$$

and

$$p^*_0 = p^M + \frac{(1-\gamma)(6\mu^2+\gamma(-1+2\mu+9\mu^2))}{2(1+2\gamma)(-2+8\mu^2+\gamma(-1-2\mu+5\mu^2))}$$

If extracting rent by means of fixed fees is sufficiently costly — i.e., $\mu > 1$ *— then* $p^*_d < p^*_{d,0}$ *and* $p^*_1 < p^*_0$.

Hence, when extracting profits by means of a fixed fee is costly (e.g., because of frictions akin moral hazard or adverse selection), platform parity benefits consumers even when two-part tariffs are implemented along the supply chain, as long as the cost of extracting profits up-front is sufficiently high.

Commitment in the wholesale model. In the baseline wholesale model, we have assumed that M sets the price on the direct channel at the last stage of the game (i.e.,

simultaneously with the intermediaries). By contrast, we now posit that M can commit to p_d before contracting takes place. In this new game, the intermediaries and the platforms solve the same maximization problems as in the baseline model, with the only difference that, in both parity regimes $k = 1, 0$, their choices depend on the actual value of p_d (an not on the equilibrium value $p^*_{d,k}$). Hence, M's choice at $t = 1$ internalizes its impact on the (optimal) wholesale and retail prices set along the supply chain.
The (symmetric) equilibria of the two regimes are summarized and compared in the following result.

Proposition 2.C.2. *Suppose that M can credibly commit to the price on the direct channel before contracting. Then:*

(i) M never distorts the price on the direct channel — i.e., $p^*_{d,1} = p^*_{d,0} = p^M$.

(ii) Retail prices on the indirect channel are distorted upward compared to p^M — i.e.,

$$p^*_1 = p^M + \underbrace{\frac{3(1-\gamma)}{2(4+\gamma)}}_{\textit{Multiple mark-ups}}.$$

and

$$p^*_0 = p^M + \underbrace{\frac{3(2-\gamma-\gamma^2)}{2(8+5\gamma)}}_{\textit{Multiple mark-ups}}.$$

with $p^*_1 < p^*_0$. Hence, platform parity still benefits consumers.

Hence, when M can credibly commit to the price charged in the direct channel, it acts as a Stackelberg leader and sets that price efficiently regardless of whether platform parity is in place or not. Therefore, the beneficial effect of platform parity is the same as in the baseline model, except that it must be diluted from the channel externality which wipes out because of commitment. The same results hold true when considering asymmetric (horizontal) product differentiation. Details are available upon request.

$N > 2$ **competing platforms in the wholesale model.** We now investigate the effect of increased competition in the indirect channel on the welfare effects of platform parity in the wholesale model. To do so, we assume that M deals with $N \geq 2$ symmetric platforms, each of them being in exclusive relationship with one intermediary. Finally, as in the (agency) model of Section 2.3.6, we assume that platform parity either applies to all platforms or to none, and we consider the following (direct) demand functions:

$$q_i(\cdot) = \frac{1-\gamma}{1+(N-1)\gamma - N\gamma^2}\left(1 - \frac{1+(N-1)\gamma}{1-\gamma}p_i + \frac{\gamma}{1-\gamma}\sum_{j\in\mathcal{I}\setminus\{i\}} p_j\right),$$

with $\mathcal{I} \triangleq \{1,\ldots,N,d\}$.

We can show the following result.

Proposition 2.C.3. *With platform parity final prices are*

$$p^*_{d,1} = p^M + \frac{3\gamma N}{4\left(4+(N-1)\gamma\right)(1+N\gamma)}, \qquad p^*_1 = p^*_{d,1} + \frac{16(1-\gamma)+N\gamma(7-8\gamma)}{12(1+N\gamma)\left(4+(N-1)\gamma\right)}.$$

By contrast, in the absence of platform parity, final prices are

$$p^*_{d,0} = p^M + \frac{3N\gamma\left(1+(N-1)\gamma\right)}{2\left(8+(5N-1)\gamma\right)(1+N\gamma)}, \qquad p^*_0 = p^*_{d,0} + \frac{3(1-\gamma)\left(2+(2N-1)\gamma\right)}{2(1+N\gamma)\left(8+(5N-1)\gamma\right)}.$$

Moreover, $p^*_{d,1} < p^*_{d,0}$ and $p^*_1 < p^*_0$. Hence, platform parity benefits consumers.

Figure 2.18 shows that the beneficial effect of platform parity is more pronounced in more competitive industries since the multiple marginalization problem becomes less relevant as the indirect channel becomes less concentrated — i.e., as N grows large. To be more specific, notice that the difference between the pass on rate with and without parity,

$$\frac{\partial p^*_1(t)}{\partial t} - \frac{\partial p^*_0(t_i)}{\partial t_i} = \frac{1+\gamma(N-1)}{4+\gamma(N-1)} - \frac{1}{4},$$

is increasing in N: as the market becomes more competitive, intermediaries are more sensitive to cost variations. As a result, the beneficial effect of platform parity is more pronounced as competition in the industry intensifies.

Alternative bargaining processes in the agency model. We now show that alternative assumptions concerning the bargaining process over platforms' commissions in the agency model strengthen the pro-competitive effect of a platform parity agreement. Specifically, we consider two bargaining process.

First, we assume that M bilaterally bargains with each platform P_i over the commission f_i and, eventually, a platform parity agreement is included in the contract. Similarly to Rey and Vergé (2017), we consider *efficient contracting* — i.e., we assume that, in each negotiation, M and P_i agree on a value f_i which maximizes their joint profit, and the surplus generated by the successful bargaining is then shared between the two players

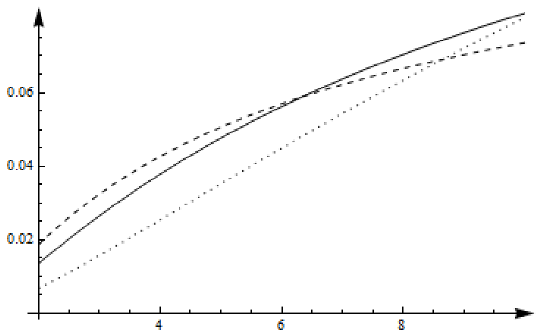

Fig. 2.18: Difference between the consumer surplus in the presence and in the absence of a platform parity agreement, in the wholesale model with N competing platforms and intermediaries, as a function of $N \geq 2$. Parameter's value: $\gamma = 0.3$ (dotted line), $\gamma = 0.6$ (continuous line), $\gamma = 0.9$ (dashed line).

according to exogenously given weights.[29]
Alternatively, we consider the polar opposite of the bargaining process examined in the paper, namely the case in which M has all the bargaining power *vis-à-vis* the platforms. In both cases, at later stages, the game unravels as in the baseline agency model analyzed in the paper.
Our results are summarized by the following Proposition.

Proposition 2.C.4. *If M either makes a take-it-or-leave-it offer f_i to each platform, or bilaterally bargains with each platform over f_i to maximize their joint profit, then, regardless of the presence of a platform parity agreement, these commissions are set to zero. Accordingly, in the two regimes, the monopolist sets access prices*

$$\tau_0^* = \frac{(2+3\gamma)^2}{8+22\gamma+12\gamma^2} > \tau_1^* = \frac{2+5\gamma+3\gamma^2}{4+10\gamma+4\gamma^2}$$

on both platforms, whereas prices on the direct channel are

$$p_{d,0}^* = p^M + \frac{\gamma(1+\gamma)}{4+11\gamma+6\gamma^2} > p_{d,1}^* = p^M + \frac{\gamma}{4+10\gamma+4\gamma^2}.$$

Finally, both intermediaries set retail prices

$$p_0^* = \tau_0^* + \frac{1-\gamma}{4+3\gamma} > p_1^* = \tau_1^* + \frac{1-\gamma}{4+2\gamma}.$$

[29]Therefore, to determine firms' equilibrium profits, we should specify how the *disagreement payoffs* are computed. However, the impact of a platform parity agreement on consumer surplus only depends on the equilibrium value of the commissions, which in turn, in the considered bargaining scheme, will turn out to be independent on the *disagreement payoffs*.

Hence, the adoption of a platform parity agreement is always beneficial to consumers.

Thus, when platform parity does not lead to an increase in platforms' commissions, it unambiguously lowers access prices and, in turn, final prices.

Notice that, since, in both regimes $k = 1, 0$, platforms get per-unit commissions $f_k^* = 0$, the considered equilibrium is equivalent to the equilibrium outcome of an industry with only one level of intermediation, in which the wholesale model is adopted — i.e., a market in which the monopolist sets the wholesale prices τ_i to platforms, which compete in the final market by setting final prices p_i.

Finally, in the framework at hand, platform parity is always beneficial to consumers also when considering asymmetric (horizontal) product differentiation. Details are available upon request.

Chapter 3

On the Ratchet Effect in Competitive Markets

3.1 Introduction

When the contractual relationship between a principal and an agent is repeated over time and the principal lacks commitment power (spot contracting), the *ratchet effect* considerably complicates the structure of the optimal dynamic contract when the agent's type is constant (or correlated) over time (see, e.g., Freixas et al. 1985, Laffont and Tirole, 1987, 1988, and Hart and Tirole, 1988, among many others). As the relationship evolves, information about the agent's type is revealed and, due to her inability to commit to a long-term contract, the principal cannot refrain from using this information in order to capture the agent's surplus later on in the relationship. Anticipating this, the agent will alter his early behavior in order to secure future rents. For this reason, dynamic incentives may be much more complex than static ones. In particular, compared to a static environment, separation is harder to sustain early on in the relationship: information is gradually revealed over time and, as a consequence, repeated relationships typically feature greater efficiency in later stages.

Non-commitment situations of this kind are very common in long run contractual relationships between private parties, and in particular in organizations. Contracts are costly to write and contingencies are often hard to foresee, which gives rise to the allocation of discretion to some members of the organization (as emphasized by Coase, 1937, Simon, 1951, Williamson, 1973, and Grossman and Hart, 1986, *inter alia*). The existence of discretion creates scope for the parties who exercise this discretion to use the information revealed by the other parties.

To the best of our knowledge, the available literature on the ratchet effect deals with these issues abstracting from the consideration of any interaction among organizations — i.e., assuming that the payoff obtained by a principal only depends on the output of the activity exerted by her agent. However, in many instances, more organizations coexist in the same market, as recognized by another well established strand of the literature, which explicitly models competition among managerial firms (although in a static framework). How does the introduction of payoff externalities, deriving from market competition, modify the analysis of dynamic incentive schemes in long-term relationships under non-commitment? The answer to this question is far from being trivial. The benefit that, in a competitive environment, a principal obtains from learning the private information of his agent at any given stage clearly depends on the expected market prices in the current and in future periods. This consideration suggests that coordination problems may emerge, which in turn creates the scope for the existence of mixed strategy equilibria, in which principals, as well as agents, in one or more periods, randomize between offering different contracts and accepting the offered contracts, respectively. Novel kind of equilibria are thus likely to arise as a consequence of the introduction of payoff externalities in a model of repeated adverse selection and short-term contracts. The implications of these novel equilibria, especially in terms of industry dynamics, may well be different as compared to those identified by the previous literature in the absence of competition, and somewhat counterintuitive, as we are going to show.

In this paper, we examine the relationship between competition and the ratchet effect. To this purpose, we study a simple two-period economy populated by a continuum of perfectly competitive firms, each composed by a principal and an exclusive agent who is privately informed about his (persistent) production cost. For simplicity, we consider a discrete characteristics space, specifically a two-type case. Following the literature (e.g., Hart and Tirole, 1988) we assume that principals lack commitment power and can only use spot contracts — i.e., in every period they simultaneously offer a wage to their agents. The interplay between competition and the ratchet effect in this simple environment has new interesting effects. To begin with, we show that, in a static environment, when the adverse selection problem takes intermediate values, the game does not feature neither a symmetric equilibrium in which all principals offer the separating contract — i.e., they shut down production by the inefficient type — nor an equilibrium in which they all offer the pooling contract — i.e., they always produce regardless of the agents' type. In this region of parameters, there exist either a symmetric mixed-strategy equilibrium in which principals randomize between offering the separating and the pooling contract, or an asymmetric, payoff equivalent, equilibrium in which a fraction of principals offers the separating contract and the rest the pooling one. The reason is that, when the market price is responsive to aggregate supply, if all principals offered the separating (pooling) contract, then the market price would be sufficiently high (low), implying that any principal would find it profitable to deviate by offering the pooling (separating) contract.

We then turn to characterize the equilibria of the two period game, assuming that players do not discount future profits. Under this assumption, the ratchet effect is magnified, since principals cannot fully screen types in the first period — i.e., information can be learned in the first period only by offering a semi-separating contract (see, e.g., Bolton and Dewatripont, 2005, Ch. 9). We show that, when the adverse selection problem is sufficiently severe, there exists a novel type of mixed strategy equilibrium in which all principals and efficient agents randomize in the first period. Specifically, principals offer the pooling contract with a given probability and the semi-separating contract otherwise. In the latter case, as in the standard analysis with inelastic demand, efficient agents randomize between accepting or not that contract. Finally, no randomizations occur in the second period, in which all efficient agents produce and obtain no rent, whereas inefficient types are shut down. The considered equilibrium has two peculiar features. First, despite learning occurs, with some probability, in the first period, when the adverse selection problem is not too severe, aggregate output is decreasing over time (*declining industry*). The reason is that, even when the adverse selection problem is sufficiently severe (such that in a static game all principals would offer the separating contract), due to the ratchet effect, to shut down production in the first period is too costly for the principals, implying that the first period is characterized by extensive pooling, whereas in the second period, in which separation of types is costless for principals, they find it optimal to let only efficient types produce. Second, the principals' expected profit increases as the adverse selection problem worsens: as the rent to be paid to the efficient type grows larger, principals shut down production in the first period with a higher probability, which results in a higher price, and in turn in a higher profit for principals.

Finally, we analyze the dynamic game under the assumption that players discount future profits, which implies that full separation is attainable in the first period. We show that, for intermediate values of the adverse selection problem, namely in the region of parameters in which the static game features the above outlined mixed strategy equilibrium, the two-period game features a novel type of semi-separating equilibrium in which principals, rather than agents, randomize between offering the pooling and the separating contracts in both periods. Interestingly, in this equilibrium, aggregate production is the same in both periods. Therefore, even tough principals can fully screen types in the first period, and learning indeed occurs with some probability in the first period, the standard result that with spot contracting efficiency increases over time may fail to hold in competitive markets. Moreover, once again, due to the presence of price externalities, principals benefit from facing a more severe adverse selection problem.

To sum up, our simple model shows that the presence of repeated market interactions among vertical hierarchies is likely to play a major role in shaping the optimal dynamic incentive schemes, in a setting with long-term principal-agent relationships ruled by short-term contracts. Specifically, once payoff externalities are taken into account, efficiency can decrease over time (i.e., *declining industries* may be observed) and principals, as well as agents, can benefit as the adverse selection problem worsens, which damages consumers. These results suggest that policy measures aimed at increasing or protecting competition in markets in which managerial firms interact over time and, within each firm, contracts are subject to renegotiation, should carefully be designed taking into due account their impact on the industry dynamics.

The rest of the article is organized as follows. After discussing the related literature, we describe the model in Section 3.2. Section 3.3 characterizes the equilibria in a static environment. Section 3.4 shows and analyzes the equilibria of the two-period game without discounting and presents the main results. In Section 3.5, we extend our analysis to the case in which players discount profits. Section 3.6 concludes. Proofs are in the Appendix.

Related Literature. Our paper, analyzing a model of dynamic contracting with competing managerial firms, is related to two different strands of the literature. On the one hand, we contribute to the literature on competing vertical managerial firms, in which, to the best of our knowledge, only static models have been considered so far. Within this literature vein, several contributions examine price and/or quantity competition in oligopolistic markets with managerial firms, each one composed by a profit oriented owner (principal) and a self interested manager (agent), in models with hidden information and/or hidden action: see, among the earliest contributions, Caillaud et al. (1995), Martin (1993), Raith (2003). More closely similar to our framework, Hart (1983) and Scharfstein (1988), dealing with managerial slack, were among the first to consider competing managerial firms in a perfect competitive market. Actually, in these works, two different types of firms, namely (profit-maximizers) entrepreneurial firms and managerial firmsa, coexist in the market. A model of hidden information with a continuum of managerial firms competing in a perfect competitive market has been recently proposed by Kastl et al. (2018), who contribute to the literature on selling information. By considering a two-period model, we add to this literature relevant insights concerning industry dynamics.
On the other hand, our paper is related to the literature on dynamic adverse selection. More specifically, we contribute to the literature on repeated adverse selection with limited commitment, in dynamic contracting situations where the agent's type is fixed over time. The seminal contributions by Stokey (1981), Fudenberg and Tirole (1983), Gul et al. (1986) focus on a durable good monopoly, showing that Coasian dynamics (Coase, 1972)

arises under seller's limited commitment[1]. The durable good model is compared to the rental (or non-durable good) model in Hart and Tirole (1988): for both cases, the authors derive the solution under full-commitment (in which the parties can commit themselves to a mechanism or contract once and for all), the solution with long-term contract and renegotiation (in which the parties can write a long-term contract, but cannot commit themselves not to renegotiate this contract by mutual agreement),[2] and the solution under limited commitment, in which the parties can only write short-term contracts which rule within a period. In the latter case, which is the one considered in our model, the ratchet effect arises, as it was already shown in two-period principal-agent models by Freixas et al. (1985) and Laffont and Tirole (1987, 1988), which bear a certain resemblance to the renting framework in Hart and Tirole (1998). As we already pointed out, all these works examine a principal-agent relationship taken in isolation, whereas we explicitly model market interactions among vertical hierarchies.

Finally, since we examine the impact of competition on the ratchet effect, our work is also related to papers investigating several others circumstances which may mitigate the ratchet effect: see, e.g., Kanemoto and MacLeod (1992), in which an external source of contract enforcement, namely competition for second-hand workers, is introduced; Carmichael and MacLeod (2000), who allow for non-Markovian strategies, by considering the threat of future punishment to sustain cooperation between an infinitely lived firm and a stream of short lived workers;[3] Acharya and Ortner (2017), who introduce in the framework of Hart and Tirole (1998) productivity shocks in the principal's benefit from having the agent work; Beccuti and Möller (2018), who, in a general mechanism design setting, analyze a model in which the seller is more patient than the buyer; Gerardi and Maestri (2018), who, assuming that a principal offers short-term contracts to an agent who can end the relationship in any period, show that rehiring is another possible remedy to the ratchet effect.

[1]For more recent contributions on the durable good monopoly under limited commitment in a mechanism design framework, with one or more potential buyers, see Skreta (2006, 2015). See also Ortner (2017), who considers a continuum of buyers and stochastic production costs.

[2]See also Laffont and Tirole (1990), who characterize the equilibrium of a two-period procurement model with commitment and renegotiation. The renegotiation issue in an infinite horizon framework has been recently addressed by Strulovici (2017) and Maestri (2017).

[3]For earlier contributions on the ratchet effect in models of labor contracting, see, e.g., Gibbons (1987), Dewatripont (1989) and Hosios and Peters (1993).

3.2 The model

Market and Players. Consider a two-period economy in which, in every period $t = 1, 2$, there is a perfectly competitive market with a of unit mass of risk-neutral firms that produce a homogeneous good. In every period there is a representative consumer with a smooth quasi-linear utility function

$$u(x_t) - p_t x_t,$$

where $x_t \geq 0$ represents the quantity consumed and p_t the market price, with $u'(\cdot) > 0$ and $u''(\cdot) \leq 0$. Since consumers take the price p_t as given, the first order condition for utility maximization, $u'(x_t) = p_t$, yields a standard differentiable downward-sloping demand function $D(p_t) \triangleq u'^{-1}(p_t)$.

In every period, firms also take the (correctly anticipated) market price p_t as given when choosing how much to produce. Each firm owner (principal, P_i) relies on a self-interested and risk-neutral manager (agent, A_i) to run the firm. In every period a firm i's production technology depends on the agent's (private) marginal cost of production $\theta_i \in \Theta \triangleq \{\underline{\theta}, \overline{\theta}\}$, with $\Delta \triangleq \overline{\theta} - \underline{\theta} > 0$ and $\Pr[\theta_i = \underline{\theta}] = \nu$, which realizes once and for all at the outset of the game (see, e.g., Freixas et al. 1985, Hart and Tirole, 1988, and Laffont and Tirole, 1987, among others). The parameter Δ can be interpreted as a measure of the severity of the adverse selection problem between principals and agents: the larger is Δ, the more relevant is agents' private information.

We assume that each firm either produces 1 unit of the good, or it does not produce at all — i.e., firm i's supply in period t is $y_{it} \in \{0, 1\}$. A binary production technology can be interpreted as an approximation of symmetric firms' production decisions in a perfectly competitive market, where firms are price takers and can either produce zero or a fixed share of the total quantity demanded. This is equivalent to assuming that firms are capacity constrained, so that no firm can supply the whole market, due for example to unmodelled technological constraints that prevent them from arbitrarily increasing the quantity produced (as in the shipping and transportation industries, electricity markets, and the hospitality industry).[4]

Hence, in every period t, aggregate supply is

$$y_t \triangleq \int_0^1 y_{it} di,$$

[4]Since firms may be ex post asymmetric and have different marginal cost of production, with unbounded supply a low-cost firm would always be able to reduce the market price up to the point where production is unprofitable for a high-cost firm, thus driving it out of the market.

and the market clearing condition requires

$$p_t \triangleq u'(y_t) .$$

We assume that with complete information — i.e., if there is no uncertainty about the agents' costs — it is always profitable for firms to produce, even when the cost is high (which is consistent with the adverse selection literature: see, e.g., Kastl et al., 2018). This requires that the lowest possible market price (when all firm produce) is sufficiently high.

Assumption 3.1. $u'(1) \geq \overline{\theta} \quad \iff \quad \Delta \leq u'(1) - \underline{\theta}$.

Moreover, to obtain interesting results, we introduce the following restriction.

Assumption 3.2. *The difference $u'(\nu) - u'(1)$ is not too large.*

This assumption requires that the utility function of the representative consumer is not too concave.

Payoffs and Contracts. In every period $t = 1, 2$, if a firm produces, given a price p_t and a transfer (wage) w_{it} paid by the principal to the agent, the principal obtains Bernoulli utility equal to $p_t - w_{it}$ and the agent obtains Bernoulli utility equal to $w_{it} - \theta_i$. Players' outside option is normalized to zero without loss of generality. For simplicity, in the baseline model, we assume that the discount factor (common to all players) is $\delta = 1$ — i.e., there is no discounting. This assumption magnifies the ratchet effect, which is at the core of our analysis. We consider discounting — i.e., $\delta \in (0, 1)$ — in Section 3.5.
Following Hart and Tirole (1988) we assume that there is no commitment — i.e., principals cannot commit in the first period to a wage to be offered in the second period (spot contracting) — and, without loss of generality, we posit that principals post a wage in every period — i.e., each agent i receives a wage offer w_{it} in every period t, which he can either accept ($x_{it} = 1$) or reject ($x_{it} = 0$).

Timing. The timing of the game is as follows. Agents learn their marginal costs at the outset of the game. Then, in every period $t = 1, 2$:

- Principals (simultaneously) offer wages and agents choose whether to accept them;
- Firms produce, wages are paid, and goods are traded.

Equilibrium. A symmetric equilibrium in pure strategies is such that: (i) each principal P_i offers a wage in the first period w_1^* and a wage in the second period $w_2^*(h_i)$ contingent on the history of the game $h_i \triangleq (p_1, x_{i1})$ observed up to period 1; (ii) given his cost θ_i, each agent A_i decides whether to produce or not — i.e., $x_t^*(\theta_i) : \Theta \to \{0,1\}$ for every period $t = 1, 2$; (iii) the aggregate supply function (because of the continuum of firms and the law of large numbers) is almost surely equal to[5]

$$y_t^* \triangleq \mathbb{E}\left[y_t^*(\theta)\right],$$

and (iv) the equilibrium price $p_t^* \triangleq u'(y_t^*)$ equalizes demand and aggregate supply (so that the product market clears). Mixed strategies are randomizations over pure strategies (more below).

Although we consider symmetric equilibria, we will argue that, by the law of large numbers, symmetric equilibria in mixed strategies can be interpreted as asymmetric equilibria in pure strategies in which principals offer different contracts.

3.3 The static benchmark

To gain intuition, we will first consider the static version of the game. In this case there are two possible symmetric equilibria in pure strategies. Namely, a separating equilibrium in which every principal offers $w^* = \underline{\theta}$ and only the efficient type produces, and a pooling equilibrium in which every principal offers $w^* = \overline{\theta}$ so that both types produce.

Consider first a separating equilibrium. By the law of large numbers, in this equilibrium aggregate production is $y^* = \nu$, and the market clearing price is such that $p^* = u'(\nu)$. Therefore, principals' equilibrium profit is $\nu\left(u'(\nu) - \underline{\theta}\right)$, while (because firms are price takers) a deviation to a pooling contract yields $u'(\nu) - \overline{\theta}$. Hence, the game features a separating equilibrium if and only if

$$\nu\left(u'(\nu) - \underline{\theta}\right) \geq u'(\nu) - \overline{\theta} \iff \Delta \geq \Delta_1 \triangleq (1-\nu)\left(u'(\nu) - \underline{\theta}\right).$$

Essentially, if the adverse selection problem is sufficiently severe — i.e., if the rent Δ that has to be paid by a principal to her efficient type is sufficiently large — every principal prefers to shut down the inefficient type.

Next, consider a pooling equilibrium in which every principal offers $w^* = \overline{\theta}$. By the law

[5] As in Legros and Newman (2013), the aggregate supply should be interpreted as a "short run" supply curve, when there is no exit or entry of new firms in the market.

of large numbers, in this equilibrium aggregate production is 1 and the market clearing price is such that $p^* = u'(1)$. Therefore, principals' equilibrium profit is $u'(1) - \overline{\theta}$, while (because of perfect competition) a deviation to a separating contract yields $\nu(u'(1) - \underline{\theta})$. Hence, the game features a pooling equilibrium if and only if

$$u'(1) - \overline{\theta} \geq \nu(u'(1) - \underline{\theta}) \iff \Delta \leq \Delta_0 \triangleq (1-\nu)(u'(1) - \underline{\theta}).$$

Clearly, if the adverse selection problem is not too severe, principals prefer to pay an information rent to the efficient type rather than shutting down the inefficient type.

However, because firms always produce in the pooling equilibrium, the market clearing price is lower than in a separating equilibrium. In other words, in a pooling equilibrium principals exert a (negative) price externality one on each other, which is reflected by the fact that $\Delta_1 > \Delta_0$. As a result, in the region of parameters where $\Delta \in (\Delta_0, \Delta_1)$ there are no symmetric equilibria in pure strategies, and if a symmetric equilibrium exists, it must involve randomizations by the principals. Assuming that every principal offers $w^* = \underline{\theta}$ with probability α^* and $w^* = \overline{\theta}$ with probability $1 - \alpha^*$, it must be

$$\underbrace{u'(\alpha^*\nu + (1-\alpha^*)) - \overline{\theta}}_{\text{Profit if } w=\overline{\theta}} = \underbrace{\nu(u'(\alpha^*\nu + (1-\alpha^*)) - \underline{\theta})}_{\text{Profit if } w=\underline{\theta}}.$$

Summing up we can state the following.

Proposition 3.1. *The equilibria of the static game are as follows.*

- *For $\Delta \leq \Delta_0$ the game features a unique symmetric equilibrium in which every principal offers $w^* = \overline{\theta}$.*
- *For $\Delta \in (\Delta_0, \Delta_1)$ there is a unique symmetric equilibrium in mixed strategy such that each principal offers $w^* = \underline{\theta}$ with probability α^* and $w^* = \overline{\theta}$ with probability $1 - \alpha^*$, with*
$$\alpha^* \triangleq \frac{1}{1-\nu}\left[1 - u'^{-1}\left(\underline{\theta} + \frac{\Delta}{1-\nu}\right)\right] \in (0,1). \tag{3.1}$$
- *For $\Delta \geq \Delta_1$ the game features a unique symmetric equilibrium in which every principal offers $w^* = \underline{\theta}$.*

The existence of the equilibrium in which principals randomize is a direct consequence of market competition. In particular, the area in which the mixed strategy equilibria exist expands as the price becomes more responsive to quantity — i.e., $\Delta_1 - \Delta_0 = (1-\nu)(u'(\nu) - u'(1))$. Since neither a pooling nor a separating equilibrium exists for intermediate values of Δ, principals must be indifferent between producing in both states or shutting down inefficient types. Hence, even if principals are ex ante identical, they may choose different contracts in equilibrium. Specifically, some of them will shut down the inefficient type, but the market price compensates them for this loss.

Remark. The law of large numbers implies that, when the game features an equilibrium in mixed strategies, it also features an asymmetric equilibrium in pure strategies such that a measure α^* of principals offers $w^* = \underline{\theta}$ and a measure $1 - \alpha^*$ offers $w^* = \overline{\theta}$. This equilibrium is in fact payoff equivalent to that in mixed strategies. By construction, if a fraction α^* of principals offer $w^* = \underline{\theta}$ while a fraction $1 - \alpha^*$ offers $w^* = \overline{\theta}$, each principal is indifferent between offering the separating and the pooling contract. This implies that compared to the case of inelastic demand — i.e., where $u''(\cdot) = 0$ and $\Delta_1 = \Delta_0$ — competition in the product market may induce ex ante identical principals to choose different contracts.

3.4 Equilibrium with repeated interaction

We now examine the two-period model and study how the interplay between competition in the product market and limited commitment alters the principals' equilibrium behavior with respect to the static benchmark. In doing so, to simplify exposition, we divide the space of parameters in three relevant regions:

1. Weak adverse selection: $\Delta \leq \Delta_0$;
2. Moderate adverse selection: $\Delta \in (\Delta_0, \Delta_1)$;
3. Strong adverse selection: $\Delta \geq \Delta_1$.

Before characterizing the equilibrium in each region, it is useful to observe that, regardless of the severity of the adverse selection problem, principals cannot fully screen agents in the first period. Essentially, by assuming $\delta = 1$ we are actually considering the case in which the ratchet effect is most severe. This is because current and future profits are weighted in the same way in the agents' inter-temporal utility function, hence an efficient agent has the largest incentive to hide his type in the first period in order to gain a rent in the second period. To formalize the argument, suppose that a separating equilibrium in the first period exists. In this candidate equilibrium only the efficient type produces in the first period and obtains a wage $w_1 \geq \underline{\theta}$, and $w_2 = \underline{\theta}$ if $x_1 = 1$ and $w_2 = \overline{\theta}$ otherwise.

Therefore, if a separating equilibrium exists, the efficient type must prefer to produce in the first period and obtain no surplus in the second period rather than producing only in the second period and obtain a rent Δ — i.e., the first period wage must be such that

$$w_1 - \underline{\theta} \geq \Delta \iff w_1 \geq \overline{\theta}.$$

But the inefficient type would then also accept w_1, whereby making separation in the first period impossible. Hence, the standard equilibrium in which principals screen agents in the first period in order to fully extract their surplus in the second period does not exist. We can now turn to characterize the equilibrium set of the game in every region of parameters.

Weak adverse selection. As in the static benchmark, when $\Delta \leq \Delta_0$ also the repeated game features a unique pooling equilibrium. The intuition is straightforward: for Δ sufficiently small, principals prefer to pay the information rent to the efficient type rather than shutting down the inefficient type. Hence, any candidate equilibrium that involves separation in the first or second period cannot exist because a deviation to a pooling contract is always profitable.

Proposition 3.2. *With weak adverse selection, the dynamic game features a unique symmetric pooling equilibrium in which $w_1^* = w_2^* = \overline{\theta}$.*

To logic behind this result is as follows. When $\Delta \leq \Delta_0$ the adverse selection problem is not too severe to induce the principals to shut down production in either of the two periods. Hence, in equilibrium both types of agents produce in both periods. Notice that such pooling equilibrium maximizes static and dynamic efficiency.

Moderate adverse selection. Suppose now that $\Delta \in (\Delta_0, \Delta_1)$. First, notice that, in this region of parameters, a symmetric semi-separating equilibrium à la Hart and Tirole (1988), such that in the first period only the efficient types produces with probability γ (see, e.g., Bolton and Dewatripont, 2005, Ch. 9), does not exist. To see why, recall that the posterior probability that an agent is efficient in the second period given that he has rejected the first period offer is

$$\Pr\left[\theta = \underline{\theta} | x_1 = 0\right] = \frac{\nu\left(1-\gamma\right)}{\nu\left(1-\gamma\right)+1-\nu},$$

which is clearly decreasing in γ: the higher the probability that the efficient type produces in the first period, the lower is the probability that an agent who has not produced in the first period is efficient. In the Appendix we show that, regardless of the equilibrium prices, in this equilibrium it is optimal to set $w_1^* = w_2^* = \underline{\theta}$, so that inefficient types are always shut down. The efficient type always produces in the second period and accepts the first period with probability γ^* being defined as the probability that makes a principal indifferent between offering $w_2 = \overline{\theta}$ and $w_2 = \underline{\theta}$ in the second period — i.e., since $p_2^* = u'(\nu)$,

$$\underbrace{u'(\nu) - \underline{\theta} - \Delta}_{\text{Profit if } w_2 = \overline{\theta}} = \underbrace{\frac{\nu(1-\gamma^*)}{\nu(1-\gamma^*)+1-\nu}(u'(\nu) - \underline{\theta})}_{\text{Expected profit if } w_2 = \underline{\theta} \text{ given } x_1 = 0} \iff \gamma^* = \frac{\Delta - \Delta_1}{\nu\Delta}. \tag{3.2}$$

Notice that this probability is increasing in Δ: as the adverse selection problem becomes more severe, the principals' profit from offering $w_2 = \overline{\theta}$ drops. Hence, to satisfy the indifference condition (3.2), γ^* must increase in order to reduce the profit associated with an offer $w_2 = \underline{\theta}$. Hence, when Δ is not large enough (i.e., $\Delta \leq \Delta_1$) a semi-separating equilibrium à la Hart and Tirole (1988) does not exist. This is because with moderate adverse selection a principal has no incentive to shut down the inefficient type in the second period — i.e., it is optimal to offer $w_2 = \overline{\theta}$ for every $\gamma \in [0,1]$ when $p_2 = u'(\nu)$ (see the Appendix). Anticipating this, the efficient type is never willing to produce in the first period, which in turn implies that an equilibrium in which all principals offer $w_1 = \underline{\theta}$ cannot exist.

Next, consider a pooling equilibrium in which all principals offer $w_1^* = \overline{\theta}$. In this case, since nothing is learned in the first period, the second period is *de facto* identical to the static game analyzed in Section 3.3. Therefore, neither $w_2^* = \overline{\theta}$ nor $w_2^* = \underline{\theta}$ can be an equilibrium because the static game does not feature neither a pooling nor a separating equilibrium for $\Delta \in (\Delta_0, \Delta_1)$. Hence, following a first period offer $w_1^* = \overline{\theta}$, in the region of parameters under consideration, principals must randomize in the second period — i.e., they must offer $w_2^* = \underline{\theta}$ with probability α^* and $w_2^* = \overline{\theta}$ with probability $1 - \alpha^*$. To show that this is in fact an equilibrium, notice that, since principals are indifferent between pooling and separating in the second period, second period profits do not matter for a deviation to be profitable and profitable deviations can only occur in the first period. Therefore, we only need to check that a principal does not find it optimal to deviate to $w_1^D = \underline{\theta}$.

However, this deviation cannot be profitable since no agent is willing to accept this offer. To see this, first notice that inefficient types never accept $w_1^D = \underline{\theta}$, whereas the efficient type accepts the first period with probability γ^* such as to make the deviating principal indifferent between offering $w_2 = \overline{\theta}$ and $w_2 = \underline{\theta}$ in the second period — i.e., since $p_2^* =$

$\underline{\theta} + \frac{\Delta}{1-\nu}$,[6]

$$\underbrace{\frac{\Delta}{1-\nu}}_{\text{Profit if } w_2=\overline{\theta}} - \Delta = \underbrace{\frac{\nu\left(1-\gamma^D\right)}{\nu\left(1-\gamma^D\right)+1-\nu}\frac{\Delta}{1-\nu}}_{\text{Expected profit if } w_2=\underline{\theta} \text{ given } x_1=0} \iff \gamma^D = 0.$$

In words, as before, if an efficient agent were willing to accept $w_1^D = \underline{\theta}$ with any positive probability, the principal would be better off offering $w_2^D = \overline{\theta}$ in the second period, but then the agent would never accept the first period offer. Put it another way, with moderate adverse selection, a deviating principal cannot induce a semi-separating outcome. Therefore, her first period profit would be zero, whereas the second period profit would be equal to the equilibrium one, implying that the deviation profit is lower than the one she would get in the candidate equilibrium under consideration.[7]

We can thus state the following.

Proposition 3.3. *With moderate adverse selection the dynamic game features a unique symmetric equilibrium in which principals offer the pooling contract in the first period — i.e., $w_1^* = \overline{\theta}$ — and randomize in the second period — i.e., $w_2^* = \underline{\theta}$ with probability α^*, and $w_2^* = \overline{\theta}$ otherwise.*

Hence, following the logic of Proposition 3.1, in the region of parameters where Δ takes intermediate values, if a symmetric equilibrium exists, then principals must randomize. The intuition behind this proposition is simple. Since principals can partially screen agents in the first period only shutting down production with a high probability, with moderate adverse selection they all prefer to give up a rent to efficient agents. Hence, shut down of inefficient types occurs, with some probability, only in the second period, which, being identical to the static game, features the mixed strategy equilibrium characterized in Proposition 3.1.

Strong adverse selection. Finally, suppose that $\Delta \geq \Delta_1$. As seen before, in this region of parameters a semi-separating equilibrium à la Hart and Tirole (1988) might exist, since the value for γ^*, given by (3.2), is strictly positive.[8] However, other types of equilibria

[6] Clearly, the efficient agent must be indifferent between accepting or not the wage $w_1^D = \underline{\theta}$ — i.e., it must be $w_2^D = \underline{\theta}$. If on the contrary, the agent believes that, with some positive probability, the deviating principal will offer $w_2^D = \overline{\theta}$, then it would never accept $w_1^D = \underline{\theta}$. Put it another way, throughout the analysis, principals can find it optimal to offer $w_1 = \underline{\theta}$ only if this offer induces semi-separation — i.e., it leads efficient types to randomize between accepting or not the first period offer and the principal then offers $w_2 = \underline{\theta}$.

[7] Since second period profits do not matter for a deviation to be profitable, from the foregoing analysis it immediately follows that a deviation involving a randomization between $w_1^D = \underline{\theta}$ and $w_1^D = \overline{\theta}$ in the first period is never optimal either.

[8] Moreover, under Assumption 3.1, $\gamma^* < 1$.

may exist. Specifically, in the next proposition we show that, as Δ grows large, shut down of production occurs in both periods. We show that there exists a novel type of mixed strategy equilibrium in which, in the first period, principals randomize between offering $w_1 = \underline{\theta}$ and $w_1 = \overline{\theta}$ and efficient types randomize between accepting and rejecting the offer $w_1 = \underline{\theta}$.

To begin with, recall that, from the analysis of the static game, for every $\Delta > \Delta_1$, if all principals offered the pooling contract in the first period ($w_1 = \overline{\theta}$), then, in the second period, they find it optimal to offer the separating contract ($w_2 = \underline{\theta}$). Alternatively, principals can offer $w_1 = \underline{\theta}$ in the first period so to induce the efficient type to randomize — i.e., the semi-separating outcome is implemented, hence principals then offer $w_2 = \underline{\theta}$ with probability one. Therefore, regardless of the contract offered by the principals in the first period, in the second period they all offer $w_2^* = \underline{\theta}$, and obtain a profit $\nu(u'(\nu) - \underline{\theta})$. Thus, to characterize the equilibria of the game for $\Delta > \Delta_1$, we only need to compare the first period profit that a principal can obtain by offering $w_1 = \overline{\theta}$ (i.e., the pooling contract) or $w_1 = \underline{\theta}$ (i.e., inducing the semi-separating outcome).
For any market price p_1, a principal finds it optimal to offer the pooling contract if and only if

$$p_1 - \overline{\theta} > \nu\gamma^*(p_1 - \underline{\theta}),$$

with γ^* being given by (3.2), yielding

$$\Delta < \sqrt{\Delta_1(p_1 - \underline{\theta})}.$$

In the Appendix, we prove that, if every other principal offers the pooling contract (i.e., $p_1 = u'(1)$), then any principal finds it optimal to offer the pooling contract as well for every $\Delta < \Delta_2$, where

$$\Delta_2 \triangleq \sqrt{\frac{\Delta_0\Delta_1}{1-\nu}}.$$

If, on the contrary, every other principal offers the semi-separating contract (i.e., $p_1 = u'(\frac{\Delta-\Delta_1}{\Delta})$), then a deviation to the pooling contract is profitable for every $\Delta < \Delta_3$, where $\Delta_3 > \Delta_2$ denotes the unique solution of

$$\frac{\Delta^2}{\Delta_1} = u'\left(1 - \frac{\Delta_1}{\Delta}\right) - \underline{\theta}. \tag{3.3}$$

As a consequence, for every $\Delta \in (\Delta_2, \Delta_3)$, a symmetric equilibrium of the game must involve principals' randomization in the first period. Specifically, we can show the following.

Proposition 3.4. *With strong adverse selection the game features the following symmetric equilibria.*

- *For $\Delta \in (\Delta_1, \Delta_2]$, there is a unique equilibrium in pure strategies such that every principal offers $w_1^* = \overline{\theta}$ in the first period and $w_2^* = \underline{\theta}$ in the second period.*
- *For $\Delta \in (\Delta_2, \Delta_3]$, there is a unique mixed strategy equilibrium such that every principal offers $w_1^* = \underline{\theta}$ with probability*

$$\rho^* = \frac{\Delta}{\Delta_1}\left[1 - u'^{-1}\left(\underline{\theta} + \frac{\Delta^2}{\Delta_1}\right)\right] \tag{3.4}$$

  and $w_1^* = \overline{\theta}$ otherwise. When $w_1^* = \underline{\theta}$ is offered, the efficient types accept the offer with probability γ^*. In the second period, every principal offers $w_2^* = \underline{\theta}$.
- *For $\Delta > \Delta_3$, there is a unique semi-separating equilibrium à la Hart and Tirole (1988) such that $w_1^* = w_2^* = \underline{\theta}$ and the efficient types accept the first period offer with probability γ^*.*

Recall that in Proposition 3.3 shut down only occurred with some probability in the second period. By contrast, with strong adverse selection principals are forced to shut down also in the first period. Indeed, for $\Delta \in (\Delta_1, \Delta_2]$ inefficient types are actually completely shut down in the second period (as it happens in the static game when Δ exceeds Δ_1) but not in the first period. The most interesting region of parameters occurs when principals start shutting down production also in the first period. In fact, when $\Delta \geq \Delta_2$ a principal has an incentive to shut down the inefficient type in the first period given that the others are offering the pooling contract in that period, and this deviation clearly destroys the equilibrium where all principals offer the pooling contract in the first period, because the market price is too low relative to Δ. Yet, since it is impossible to fully separate types in the first period, shut down can only occur by implementing the semi-separating outcome, implying that in the first period also efficient agents randomize. Of course, when Δ becomes too large (i.e., $\Delta > \Delta_3$) only the semi-separating equilibrium à la Hart and Tirole (1988) emerges: inefficient types are completely shut down in both periods and efficient types randomize between accepting and rejecting the first period offer.
As in the static model, the existence of an equilibrium in which principals randomize is a direct consequence of market competition (in fact, it can be easily checked that, with inelastic demand, $\Delta_2 = \Delta_3$), and it relates to the presence of price externalities. However, important differences arise with respect to the static benchmark, due to the impossibility of (fully) screening types in the first period. First, this implies that, in the first period, also

(efficient) agents randomize. Second, differently from the static game, in which pooling occurs only when $\Delta < \Delta_0$, the fact that, by offering $w_1 = \underline{\theta}$, also efficient types are shut down with a positive probability, makes it more attractive for every principal to offer the pooling contract in the first period. Hence, even when the adverse selection problem is relatively pronounced (i.e., for every $\Delta \in (\Delta_0, \Delta_3)$), the dynamic game features, at least to some extent, pooling in the first stage: the ratchet effect.

Dynamics of aggregate output. We are now ready to examine industry dynamics — i.e., to see whether aggregate output increases or decreases over time. This is far from being a trivial question, because, in our setting, there are two opposite forces driving the dynamics of aggregate output. On the one hand, information about the agents' (persistent) type can be revealed over time, which, *ceteris paribus*, would lead the aggregate output to increase over time. On the other hand, the ratchet effect makes it impossible to fully separate types at early stages, and costly to achieve a semi-separating outcome: as a consequence, early stages are characterized by extensive pooling, and, other things being equal, this would lead the aggregate output to be decreasing over time. Which one of these two forces dominates clearly depends on the severity of the adverse selection problem. Specifically, we can show what follows.

Proposition 3.5. *There exists a threshold $\hat{\Delta} \in (\Delta_2, \Delta_3)$ such that $y_1^* < y_2^*$ if and only if $\Delta > \hat{\Delta}$.*

When the adverse selection problem is not too severe (i.e., $\Delta \in (\Delta_0, \hat{\Delta})$)[9], we observe a *declining industry.* The reason is as follows. Even when the adverse selection problem is sufficiently severe such that in a static game all principals would offer the separating contract, the ratchet effect implies that to shut down production in the first period is too costly for the principals, relative to the rents to be paid to efficient types in a pooling contract. Therefore, as we already pointed out, the first period is characterized by extensive pooling. As a consequence, aggregate output is relatively high in the first period and, with relatively high probability, no information is learned in the first period, hence, in the second period, in which the ratchet effect does not play any role, principals are more willing to separate types (since $\Delta > \Delta_0$), which results in a relatively low aggregate output. On the contrary, when the adverse selection problem is too severe (i.e., $\Delta > \hat{\Delta}$), paying rents to efficient types is extremely costly for the principals, who thus prefer to shut down production with a high probability in the first period, which result in a relatively low aggregate output. In this case, despite the ratchet effect, information is revealed in

[9]Obviously, for $\Delta \leq \Delta_0$, all principals offer the pooling contract in both periods, hence aggregate output is constant over time.

the first period with a relatively high probability, and this leads to a higher production in the second period.

Profits and Welfare. We can finally examine principals' profits and consumer welfare, confining our attention to the more interesting equilibrium — i.e., the novel type of mixed strategy equilibrium defined for $\Delta \in (\Delta_2, \Delta_3]$. In particular, we can show the following comparative statics result.

Proposition 3.6. *For all $\Delta \in (\Delta_2, \Delta_3]$, principals' expected profit is increasing in Δ, whereas consumer welfare is decreasing in Δ.*

Counter intuitively, and differently from what is observed in the equilibria of the game without price externalities, in this novel mixed strategy equilibrium principals benefit from facing a more severe adverse selection problem. The reason is rather simple. As the rent to be paid to the efficient type grows larger, principals shut down production in the first period with a higher probability (ρ^* being increasing in Δ). The resulting higher price, in terms of principals' expected profit, outweighs the increased probability of not producing at all. However, the lower aggregate output in the first period unambiguously harms consumers.[10]

Remarks. To conclude this section, we briefly discuss some simple extensions of our analysis.

Asymmetric equilibria. Once again, by the law of large numbers, equilibria in mixed strategies can be interpreted as asymmetric equilibria where players play pure strategies. Specifically, for $\Delta \in (\Delta_2, \Delta_3]$, there exists an asymmetric pure strategy equilibrium, which is payoff equivalent to the mixed strategy equilibrium described above, in which, in the first period,

- a measure ρ^* of principals offers $w_1^* = \underline{\theta}$ and a measure $1 - \rho^*$ offers $w_1^* = \overline{\theta}$;
- among the agents who have been offered $w_1^* = \underline{\theta}$, only a measure γ^* of the efficient types accepts the proposed contract.

The same reasoning applies to the other mixed strategy equilibria characterized above.

Long-term renegotiation-proof contracts. One may wonder how the equilibrium configuration of the game changes when the possibility of commitment and Pareto-improving

[10] Clearly, the second period plays no role in this result, since $y_2^* = \nu$ does not depend on Δ.

renegotiation is admitted — i.e., when principals can commit to a long-term contract, offered at the beginning of the game and which rules the relationship in both periods, however they cannot commit not to offer a new contract at the beginning of the second period that would replace the initial contract if the agent finds it acceptable. It can be easily seen that, in the game without discounting, for every outcome (i.e., market prices and inter-temporal rents and profits) that can be achieved by means of short-term contracts there exists a long-term renegotiation-proof contract which yields the same outcome. However, there exists a long-term renegotiation proof contract which allows principals to implement a fully separating outcome in the first period, which, in the game with no discounting, cannot be achieved by using spot contracts. In fact, principals can commit to a contract specifying that $w_1 = \underline{\theta}$ and $w_2 = \overline{\theta}$ for every $x_1 \in \{0, 1\}$, which is clearly renegotiation-proof. As a consequence, the equilibrium set of the game under long-term renegotiation-proof contracts differs from the one characterized above. Specifically, it can be proved that, when the adverse selection problem takes extreme values (namely, for $\Delta \leq \Delta_0$ and $\Delta \geq \hat{\Delta}$), the commitment to long-term renegotiation-proof contracts does not alter the equilibrium set of the game, whereas for intermediate values of the adverse selection problem (i.e., when $\Delta \in (\Delta_0, \hat{\Delta})$) the game under commitment and renegotiation admits three novel mixed strategy equilibria in which the long-term contract which implements the separating outcome of the model with spot contracts is offered with positive probability. In all these equilibria, the aggregate quantity is the same in the two periods and, in two of them, which only exist if demand is elastic with respect to quantity, principals' profit is increasing in Δ. All details are provided in the Appendix. Thus, the main insights of our model, namely that in competitive markets efficiency does not necessarily improve over time and principals' profit can be non-monotone with respect to the severity of the adverse selection problem, remain valid under this different specification of the contracts' space.

Concavity of the utility function. One may wonder how our results would change if the utility function $u(\cdot)$ does not satisfy Assumption 3.2 — i.e., it is sufficiently concave to satisfy

$$\frac{\Delta_1}{\Delta_0} > \frac{1}{1-\nu}.$$

Since, in these cases, Δ_1 (hence, a fortiori, Δ_2 and Δ_3) does not satisfy Assumption 3.1, it follows that only the cases with weak and moderate adverse selection survive when Assumption 3.2 does not hold. It can be easily seen that the results shown in Propositions 3.2 and 3.3 apply regardless of the concavity of the utility function.

3.5 Discounting

We now study the interplay between market competition and dynamics under the assumption that players discount second period profits — i.e., we consider a discount factor $\delta \in (0, 1)$. Within this framework, principals can achieve full separation in the first period, by offering a contract $w_1 = \underline{\theta} + \delta\Delta < \overline{\theta}$, and then extract agents' rents in the second period (see, e.g., Bolton and Dewatripont, 2005, Ch. 9).

We first argue that, in this case, the most interesting region of parameters is the one characterized by moderate adverse selection — i.e., $\Delta \in (\Delta_0, \Delta_1)$ — since, as we have already proved in the case with no discounting, under our assumptions,[11] in this region the standard equilibria à la Hart and Tirole (1988) — i.e., the pooling, separating and semi-separating equilibria — fail to exist altogether. Moreover, within this region of parameters, we characterize two novel mixed-strategy equilibria in which principals rather than agents randomize in the first and, in some cases, also in the second period.

Accordingly, from now on suppose that $\Delta \in (\Delta_0, \Delta_1)$. Consider first a separating equilibrium in which all principals offer $w_1^* = \underline{\theta} + \delta\Delta$ in the first period and then fully extract the agents' surplus in the second period. It should be clear that, if principals are unable to sustain a separating outcome in the static game (i.e., when $\Delta < \Delta_1$), *a fortiori* such an equilibrium must fail to exist in the dynamic game. The intuition is as follows: a deviation to a pooling contract $w_1^D = \overline{\theta}$ allows a principal to gain in the first period not only because the separating behavior of her rivals induces a market price which is larger than the rent obtained by the efficient type in the first period, but also because, by doing so, the deviating principal saves on the inter-temporal rent that has to be granted to an efficient type in order to induce separation.[12,13]

Next, consider a symmetric semi-separating equilibrium à la Hart and Tirole (1988) in which $w_1^* = w_2^* = \underline{\theta}$ so that inefficient types are always shut down, while efficient types randomly accept the first period offer. Since, in this candidate equilibrium, efficient agents would accept the first period offer with probability γ^* defined by (3.2), it is clear that this

[11] Specifically, besides Assumptions 3.1 and 3.2 (which now requires that $\frac{\Delta_1}{\Delta_0} < \frac{1}{\sqrt{1-\nu}}$), we also assume $u''' (\cdot) \leq 0$ — i.e., that the elasticity of demand with respect to price is increasing in the market price, or alternatively, that the demand function is concave. However, all our results also hold if $u''' (\cdot)$ is positive but not too large.

[12] Recall that efficient types anticipate that informed principals achieve full surplus extraction in the second period if they accept a separating contract in the first period. Hence, the rent given up to these types in the dynamic version of the game is larger than the static rent.

[13] Notice that, other things being equal, the deviating principal incurs a second period (efficiency) loss since in the region of parameters under consideration the inefficient type is shut down. However, this loss is second order relative to the first period gain(s) since the market price in the second period is kept low by the fact that rivals are informed about their agents' types and always produce.

equilibrium cannot exist for $\Delta < \Delta_1$ (since, in this case, it would be $\gamma^* < 0$). The reason is as in the baseline model with no discounting. When Δ is not too large, a deviation to a pooling contract in the second period is always profitable. Formally, this means that there is no mixed strategy that efficient types can play in the first period that makes principals indifferent between offering a pooling and separating contract in the second period. Essentially, any agents' randomization in the first period tends to increase so much the equilibrium price in the second period that principals find it no longer profitable to shut down the inefficient type. In the proof of Proposition 3.7 we show that, with moderate adverse selection, this type of logic also rules out equilibria in which $w_1 = \underline{\theta}$ is offered with positive probability.

Finally, consider an equilibrium in which all principals offer $w_1^* = \overline{\theta}$. Because firms are perfectly competitive — i.e., the contract offered by a single principal does not affect the market price — it should be clear that, if the pooling equilibrium does not exist in the static game, it does not exist in the dynamic game either as long as $\Delta \in (\Delta_0, \Delta_1)$. This is because, in the region of parameters under consideration, as we have seen in the game without discounting, principals randomize in the second period if they all offered the pooling contract in the first period. Hence, they are indifferent between the pooling and the separating contract in the second period. As a consequence, second period profits do not matter for a deviation to be profitable. Therefore, from the analysis of the static game, we know that each principal has a profitable deviation to a separating contract in the first period. This in turn implies that the equilibrium characterized in this region of parameters for the game with no discounting does not survive when principals can perfectly separate types in the first period and achieve full surplus extraction in the second period.

Summing up, we can state the following.

Proposition 3.7. *For $\Delta \in (\Delta_0, \Delta_1)$ the dynamic game features neither symmetric equilibria in which principals play pure strategies in the first period, nor a mixed strategy equilibrium in which agents randomize — i.e., a semi-separating equilibrium à la Hart and Tirole (1988) does not exist.*

Hence, also when payoff are discounted, following the logic of Proposition 3.1, in the region of parameters characterized by moderate adverse selection, if a symmetric equilibrium exists, then principals must randomize. Specifically, randomization occurs at least in the first period. In what follows we characterize these equilibria and analyze their properties.[14]

Clearly, since the game is repeated, randomizations may occur also in the second period.

[14] Once again, by the law of large numbers, the symmetric mixed strategy equilibria that we are going

We shall thus consider mixed strategy equilibria in which, in the first period, each principal randomizes between offering the separating and the pooling contract; in the second period, the principals who offered the separating contract learn the agent's type and fully extract his rent, while those who offered the pooling contract can, alternatively: (i) offer a separating contract; (ii) offer a pooling contract; or (iii) randomize again. Clearly, which of these options will prevail in equilibrium depends on the severity of the adverse selection problem.

Let

$$\Delta^* \triangleq (1-\nu)(u'(\tfrac{1+\nu}{2}) - \underline{\theta}),$$

with $\Delta^* \in (\Delta_0, \Delta_1)$ since $u'(\nu) > u'(\frac{1+\nu}{2}) > u'(1)$. We can show the following.

Proposition 3.8. *For $\Delta \in (\Delta_0, \Delta_1)$ a symmetric equilibrium must induce principals to randomize at least in the first period. Specifically:*

- *If $\Delta \in (\Delta_0, \Delta^*]$ there exists only one symmetric equilibrium in which principals randomize in both periods:*
  - *In the first period, each principal offers $w_1^* = \underline{\theta} + \delta\Delta$ with probability α^* and $w_1^* = \overline{\theta}$ with probability $1-\alpha^*$, with $\alpha^* \in (0, 0.5)$ being defined in (3.1).*
  - *In the second period, if she has offered $w_1^* = \underline{\theta} + \delta\Delta$, a principal offers $w_2^* = \underline{\theta}$ if $x_1 = 1$, and $w_2^* = \overline{\theta}$ otherwise. If $w_1^* = \overline{\theta}$, then $w_2^* = \underline{\theta}$ is offered with probability*

$$\varepsilon^* \triangleq \frac{1 - u'^{-1}\left(\underline{\theta} + \frac{\Delta}{1-\nu}\right)}{u'^{-1}\left(\underline{\theta} + \frac{\Delta}{1-\nu}\right) - \nu}, \tag{3.5}$$

    and $w_2^* = \overline{\theta}$ with probability $1-\varepsilon^*$. Moreover, $\varepsilon^* > \alpha^*$.

- *If $\Delta \in (\Delta^*, \Delta_1)$ there exists only one symmetric equilibrium in which principals randomize in the first period only. That is:*
  - *In the first period each principal offers $w_1^* = \underline{\theta} + \delta\Delta$ with probability β^* and $w_1^* = \overline{\theta}$ with probability $1-\beta^*$, with $\beta^* \in [0.5, 1]$ being the unique solution of*

$$\frac{1-\delta}{1-\nu}\Delta = u'(\nu\beta^* + 1 - \beta^*) - \underline{\theta} - \delta(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta}). \tag{3.6}$$

  - *In the second period, each principal offers $w_2^* = \underline{\theta}$ if in the first period she has offered $w_1^* = \overline{\theta}$ or if she has offered $w_1^* = \underline{\theta}+\delta\Delta$ and $x_1 = 1$. Otherwise, $w_2^* = \overline{\theta}$ is offered.*

to derive can be alternatively interpreted as asymmetric pure strategy equilibria. Details are available upon request.

With an elastic demand and discounting, two novel equilibria arise in the region of parameters characterized by moderate adverse selection. Indeed, these two equilibria cannot exist when demand is inelastic (since principals never randomize in this case), nor when future profits are not discounted (since principals cannot achieve full separation in the first period).
In the first equilibrium, which exists when Δ is not too large (i.e., $\Delta \leq \Delta^*$), principals mix in both periods and randomizations are correlated over time, since, for obvious reasons, a principal randomizes in the second period only when she has not learned the agent's type in the first period (i.e., when the outcome of the randomization in the first period was a pooling contract). The reason why this equilibrium exists for $\Delta \leq \Delta^*$ is as follows: when Δ is relatively small, granting rents to the efficient type is not too costly for a principal, so that if all other principals shut down the inefficient type, there is an incentive to offer the pooling contract since the market price is high relative to the severity of the adverse selection problem. By contrast, if every principal pools, then offering a separating contract is optimal from an individual point of view because in the pooling equilibrium the market price is too low relative to Δ (recall that $\Delta > \Delta_0$).
When Δ is large enough (i.e., $\Delta > \Delta^*$) principals randomize in the first period only, and the probability β^* of selecting a separating contract must be such that they are indifferent between offering $w_1^* = \underline{\theta} + \delta\Delta$ and $w_1^* = \overline{\theta}$. Hence, the indifference condition that pins down β^* equalizes the inter-temporal profits of each principal when she offers a pooling and a separating contract. In the second period each principal fully extracts the agent's surplus if she has learned his type, otherwise the separating contract is offered. In fact, since in this area Δ is sufficiently large, an uninformed principal always prefers to shut down the inefficient type in order to save on the rent given up to the efficient type. Hence, there cannot be played mixed strategies in the second period.
Clearly, $\beta^* > \alpha^*$ since, as the adverse selection problem becomes more pronounced, principals prefer to shut down the inefficient types with higher probability in order to save on the efficient type's rent.

Dynamics of aggregate output. Interestingly, when the equilibrium features randomizations in both periods, the market price is constant over time — i.e.,

$$p_1^* = p_2^* = \underline{\theta} + \frac{\Delta}{1-\nu}.$$

As a result, although randomizations occur in this equilibrium, aggregate quantity is constant and not increasing over time, as it happens in a model with inelastic demand à la Hart and Tirole (1988) when a separating or a semi-separating equilibrium emerges.

This shows that, even tough here principals can fully screen types in the first period (given that $\delta < 1$), the standard result that with spot contracting efficiency increases over time may fail to hold in competitive markets, even when learning occurs with some probability in the first period — i.e., when the adverse selection problem is relatively severe such that principals do not always offer the pooling contract. By contrast, for $\Delta > \Delta^*$, aggregate quantity is increasing over time as in the standard analysis. To sum up, we can thus state what follows.

Proposition 3.9. *If $\Delta \in (\Delta_0, \Delta^*]$ aggregate production is the same in the first and second period — i.e.,*

$$\underbrace{\nu\alpha^* + 1 - \alpha^*}_{y_1^*} = \underbrace{\alpha^* + (1-\alpha^*)\,(\nu\varepsilon^* + 1 - \varepsilon^*)}_{y_2^*}.$$

If $\Delta \in (\Delta^*, \Delta_1)$ aggregate production is higher in the second period than in the first period — i.e.,

$$\underbrace{\nu\beta^* + 1 - \beta^*}_{y_1^*} < \underbrace{\beta^* + \nu\,(1-\beta^*)}_{y_2^*}.$$

Hence, in the region of parameters where $\Delta \in (\Delta_0, \Delta^*]$ competition tends to stabilize prices and production over time. This is because, although principals who have not learned the agent's type in the first period tend to be more selective in the second period — i.e., they offer the pooling contract with a lower probability than in the first period ($\varepsilon^* > \alpha^*$) — the other principals have learned in the first period and thus always produce in the second period. The result of Proposition 3.9 shows that, in the equilibrium, the positive effect of learning exactly compensates the higher probability of shutting down production in second period.

Notably, in this equilibrium, the static outcome is repeated over time — i.e., the aggregate output in both periods is equal to the aggregate output of the static game. Roughly speaking, when the adverse selection problem is relatively weak (i.e., $\Delta \leq \Delta^*$), the possibility of learning the agent's type in the first period is not much attractive for principals, and, as a consequence, efficiency does not improve over time.

However, differently from the game without discounting, in the considered region of parameters, aggregate quantity cannot be strictly decreasing over time. Thus, when the ratchet effect is weaker, namely when full screening is viable in the first period, the separating contract is offered with positive probability in equilibrium, thereby reducing the extent of pooling in the first period and, consequently, ruling out the possibility of observing a declining industry.

Finally, it is interesting to see how introducing the possibility of full separation in the first period modifies aggregate quantities in both time periods as compared to the game without discounting.

Proposition 3.10. *For all $\Delta \in (\Delta_0, \Delta_1)$, aggregate quantity in the first period is higher when players do not discount future profits — i.e., $y_1^*|_{\delta=1} > y_1^*|_{\delta<1}$. However, in the second period,*

- $y_2^*|_{\delta=1} = y_2^*|_{\delta<1}$ *for* $\Delta \in (\Delta_0, \Delta^*]$,
- $y_2^*|_{\delta=1} < y_2^*|_{\delta<1}$ *for* $\Delta \in (\Delta^*, \Delta_1)$.

As expected, when the ratchet effect is more intense (i.e., when $\delta = 1$), the first period is characterized by extensive pooling, which entails a higher aggregate output as compared to the game in which principals can fully screen their agents' type in the first period. By the same reasoning, when Δ is sufficiently high ($\Delta > \Delta^*$), a relatively high fraction of principals finds it optimal to offer the separating contract in the first period, which yields a higher aggregate output in the second period as compared to the static game, whose outcome is achieved in the game without discounting. By contrast, when $\Delta \in (\Delta_0, \Delta_1)$, for every value of the discount factor, the aggregate output in the second period coincides with that of the static game.

Profits and Welfare. As for profits and consumer welfare in the considered equilibria, we can show the following comparative statics result.

Proposition 3.11. *For $\Delta \in (\Delta_0, \Delta^*]$, principals' equilibrium profit is increasing in Δ. This is true also for $\Delta \in (\Delta^*, \Delta_1]$ when δ is relatively small, whereas the opposite holds when δ is sufficiently high. Finally, consumer welfare is decreasing in Δ for all $\Delta \in (\Delta_0, \Delta_1]$.*

The trade-off is as follows. On the one hand, as in the mixed strategy equilibrium analyzed in the case with no discounting, when Δ increases, principals tend to shut down production with higher probability, and, *ceteris paribus*, benefit from the resulting higher prices. On the other hand, here principals can separate types in the first period only by giving up a rent $\delta\Delta$. As a matter of fact, when the adverse selection problem is sufficiently severe and δ is relatively high, this latter effect dominates, whereby leading the equilibrium profit to be decreasing in Δ.
The result concerning consumer welfare is immediate when $\Delta \in (\Delta_0, \Delta^*]$, since, for the reasons pointed out above, the aggregate output (which is constant over time) is decreasing

in Δ. As for the other equilibrium, it can be proved (see the Appendix) that

$$\frac{\partial y_1^*}{\partial \Delta} = -\frac{\partial y_2^*}{\partial \Delta} < 0,$$

which implies that, under the assumption that consumers discount future utility, consumer welfare is strictly decreasing in Δ.

Remarks. Some other results are briefly discussed in the following concluding remarks.

Weak and strong adverse selection. One may wonder what are the equilibria of the game with discounting in the cases of weak and strong adverse selection. The answer is rather simple. When $\Delta \leq \Delta_0$, it is easy to prove that, as in the game without discounting, in the unique equilibrium of the game, all principals offer the pooling contract in both periods. When $\Delta \geq \Delta_1$, within different subregions of parameters, there may exists two symmetric pure strategy equilibria, namely the separating and the semi-separating à la Hart and Tirole (1988), and the following symmetric mixed strategy equilibria, in which:[15]

- in $t = 1$, principals randomize between the semi-separating contract ($w_1 = \underline{\theta}$) and the pooling contract ($w_1 = \overline{\theta}$), and in $t = 2$ all principals offer $w_2 = \underline{\theta}$;
- in $t = 1$, principals randomize between the semi-separating contract ($w_1 = \underline{\theta}$) and the separating contract ($w_1 = \underline{\theta} + \delta\Delta$), and in $t = 2$ the principals who offered the semi-separating contract offer $w_2 = \underline{\theta}$, whereas the principals who offered the separating contract achieve full surplus extraction;
- in $t = 1$, principals randomize among the semi-separating contract ($w_1 = \underline{\theta}$), the separating contract ($w_1 = \underline{\theta} + \delta\Delta$) and the pooling contract ($w_1 = \overline{\theta}$), and in $t = 2$ the principals who offered the semi-separating and the pooling contract offer $w_2 = \underline{\theta}$, whereas the principals who offered the separating contract achieve full surplus extraction.

Long-term renegotiation-proof contracts. It can be proved that, for all $\delta \in (0, 1)$, the equilibrium outcome with spot contracts coincides with the equilibrium outcome with long-term contracts and Pareto-improving renegotiations. This result, which is well known in the two-period model without competition,[16] also holds true in our model with competition, because it can be established a one-to-one correspondence among the outcomes

[15] Some details can be found in the Appendix. Additional details are available upon request.

[16] As shown by Hart and Tirole (1988), in general, the sale and rental outcome with long-term contracts and renegotiation coincide both with one another and with the spot contracting sale outcome. Moreover,

(i.e., market prices and discounted rents and profits) that can be achieved with long-term renegotiation-proof contracts and with short-term contracts.[17]

3.6 Conclusion

The analysis of dynamic incentive schemes under limited commitment and when market competition takes place over time is a rather complex issue. In this paper, we tackled this problem by studying a two-period economy populated by a continuum of perfectly competitive firms and under the assumption of a discrete characteristics space. This simple set up enabled us to provide a complete characterization of the equilibrium set of the game.

The equilibrium analysis has revealed that, once market interaction is taken into account, some of the well known results found in the previous literature on repeated adverse selection with limited commitment and fixed agents' types do not longer hold true. Specifically, when the adverse selection problem is relatively severe, despite some principals learn their agents' type in the early stages, due to the ratchet effect, efficiency (i.e., aggregate output) may decrease over time. Moreover, a more severe adverse selection problem is likely to be translated into a higher market price, which benefits principals, while harming consumers. To explore to what extent our findings are robust along several crucial dimensions (e.g., imperfect market competition, indefinite time horizon, continuous production functions and agents' type) would be useful in order to better understand the scope for regulation in industries characterized by repeated interactions among organizations dealing with internal information asymmetries and contract renegotiation issues. We hope to address these issues in future research.

in the specific case with two periods, under non-commitment, the rental model and the sale model coincide (see, e.g., Bolton and Dewatripont, 2005, Ch. 9). Thus, the outcome of the rental model (which is equivalent to our benchmark model with inelastic demand) is the same under non-commitment and under long-term contracts and renegotiation.

[17] Thus, the mixed strategy equilibria of the game with spot contracts can be equivalently regarded as randomizations among the corresponding long-term renegotiation-proof contracts. Details are in the Appendix.

Appendix

Appendix 3.A. Proofs

Proof of Proposition 3.1. The result is proved in Section 3.3. ■

Proof of Proposition 3.2. Consider a symmetric equilibrium in which each principal offers the pooling contract in the first period ($w_1 = \overline{\theta}$). Given that, in such equilibrium, no information is extracted from the agents in the first period, the analysis of the second period is identical to the analysis of the static game, hence each principal finds it optimal to offer the pooling contract also in the second period for all $\Delta \leq \Delta_0$. In this case, her expected profit is

$$\pi^{PP} \triangleq 2(u'(1) - \overline{\theta}).$$

The considered strategy profile constitutes an equilibrium if and only if there are no profitable deviations from it — i.e., every principal must not want to deviate from the prescribed strategy profile in either period.

Since, as we already noticed, in the considered range of parameters, following a pooling contract in the first period, there are no profitable deviations in the second period, the only possibly profitable deviation for a principal consists in offering $w_1^D = \underline{\theta}$ in the first period.

However, this deviation cannot be profitable since no efficient agent is willing to accept this offer. In fact, if an efficient agent were to accept $w_1^D = \underline{\theta}$ with probability $\gamma > 0$, then the deviating principal would always find it optimal to offer $w_2^D = \overline{\theta}$ — i.e., since $p_2^* = u'(1)$, for all $\gamma > 0$ and $\Delta \leq \Delta_0$:

$$\underbrace{u'(1) - \underline{\theta} - \Delta}_{\text{Profit if } w_2 = \overline{\theta}} > \underbrace{\frac{\nu(1-\gamma)}{\nu(1-\gamma) + 1 - \nu}(u'(1) - \underline{\theta})}_{\text{Expected profit if } w_2 = \underline{\theta} \text{ given } x_1 = 0}.$$

This of course implies that, with weak adverse selection, a deviating principal cannot induce a semi-separating outcome. As a consequence, the considered deviation is not profitable, and there cannot exist other equilibria in the region of parameters under consideration. ■

Proof of Proposition 3.3. Within any managerial firm, for any given market prices p_1 and p_2, a semi-separating contract is obtained as follows. The principal offers $w_2^* = \underline{\theta}$, to an agent who has not accepted the wage $w_1 \in [\underline{\theta}, \overline{\theta})$, with a probability σ^* such that the efficient agent is indifferent between accepting or not the wage w_1. Consider the efficient type agent: if he accepts w_1, then he gets $w_1 - \underline{\theta}$ in the first period, and zero in the second period (since the principal, knowing that he is efficient, will offer $w_2 = \underline{\theta}$); if, instead, he refuses the first period contract, then he will get zero in the first period but a profit Δ with probability $1 - \sigma^*$ in the second period. Equating these payoffs yields

$$w_1 - \underline{\theta} = (1 - \sigma^*)\Delta \iff \sigma^* = 1 - \frac{w_1 - \underline{\theta}}{\Delta}.$$

Next, as detailed in Section 3.4, each efficient agent accepts w_1 with a probability γ^* such as to make the principal indifferent between offering $w_2 = \underline{\theta}$ and $w_2 = \overline{\theta}$ to an agent who refused w_1 — i.e., for any given market price p_2,

$$\underbrace{p_2 - \underline{\theta} - \Delta}_{\text{Profit if } w_2 = \overline{\theta}} = \underbrace{\frac{\nu(1-\gamma^*)}{\nu(1-\gamma^*) + 1 - \nu}(p_2 - \underline{\theta})}_{\text{Expected profit if } w_2 = \underline{\theta} \text{ given } x_1 = 0} \iff \gamma^* = \frac{1}{\nu}\left(1 - \frac{(1-\nu)(p_2 - \underline{\theta})}{\Delta}\right). \quad (3.7)$$

For any given market prices p_1 and p_2, the principal's expected profit is

$$\pi^{SS} \triangleq \nu\gamma^*(p_1 - w_1) + \pi_2,$$

where π_2 is the second-period profit, which she gets from playing according to the mixed strategy σ^* derived above. Notice that, since, by definition, π_2 must be equal to the profit that she can get by setting either $w_2 = \underline{\theta}$ or $w_2 = \overline{\theta}$ with probability one, it does not depend on w_1. Thus, π^{SS} is clearly decreasing in w_1: therefore, for every p_1, $w_1^* = \underline{\theta}$ and, accordingly, $\sigma^* = 1$ — i.e. the principal finds it optimal to offer (with probability one) $w_2^* = \underline{\theta}$ to an agent who refused w_1.
The rest of the proof is provided in Section 3.4. ■

Proof of Proposition 3.4. For $\Delta > \Delta_1$, the game can admit the following equilibria.

- First consider a pure strategy equilibrium in which all principals offer the pooling contract $w_1^* = \overline{\theta}$ in the first period. In this case, since the second period is identical to the static game, and $\Delta > \Delta_1$, all principals find it optimal to offer the separating contract $w_2^* = \underline{\theta}$ in the second period. Accordingly, in this candidate equilibrium, each principal obtains a profit

$$\pi^{PS} \triangleq u'(1) - \underline{\theta} - \Delta + \nu(u'(\nu) - \underline{\theta}).$$

As we argued in Section 3.4, a principal can only deviate from the considered strategy profile by offering $w_1^D = \underline{\theta}$, which induces the semi-separating outcome. In this case, the deviating principal obtains a profit

$$\pi^D \triangleq \left(1 - \frac{\Delta_1}{\Delta}\right) \frac{\Delta_0}{1-\nu} + \frac{\nu}{1-\nu}\Delta_1,$$

and it can be easily checked that $\pi^D < \pi^{PS}$ if and only if $\Delta < \Delta_2$.

- Next consider a mixed strategy equilibrium in which, in the first period, principals offer $w_1^* = \underline{\theta}$ (i.e., induce semi-separation) with probability ρ and $w_1^* = \overline{\theta}$ (i.e., the pooling contract) otherwise. As explained in Section 3.4, in the considered region of parameters, every principal then offers $w_2^* = \underline{\theta}$. In order for such an equilibrium to exists, principals must be indifferent between offering the two considered contracts, which immediately yields

$$\Delta^2 = (1-\nu)(u'(\rho\nu\gamma^* + 1 - \rho) - \underline{\theta})(u'(\nu) - \underline{\theta}),$$

with γ^* being defined by (3.2), from which we obtain the mixed strategy ρ^* given by (3.4). Such value turns out to be positive for $\Delta > \Delta_2$ and lower than one if Δ is *sufficiently small* to satisfy

$$\frac{\Delta^2}{\Delta_1} < u'\left(1 - \frac{\Delta_1}{\Delta}\right) - \underline{\theta} \iff \Delta < \Delta_3.$$

The result $\Delta_3 > \Delta_2$ obtains since it is easy to check that the above inequality is satisfied at $\Delta = \frac{\Delta_1}{\sqrt{1-\nu}} > \Delta_2$. Finally, notice that, since all principals find it optimal to offer $w_2^* = \underline{\theta}$ in the second period and they are indifferent between the two candidate optimal contracts in the first period, no principal can profitably deviate from this equilibrium.

- Finally, consider the symmetric semi-separating equilibrium, in which all principals offer $w_1^* = w_2^* = \underline{\theta}$, efficient agents accept the first period offer with probability γ^* defined by (3.2) and they always produce in the first period, whereas inefficient types never produce. In this candidate equilibrium, each principal would obtain

$$\pi^{SS} \triangleq \nu\gamma^* \left(u'(\nu\gamma^*) - \underline{\theta}\right) + \nu \left(u'(\nu) - \underline{\theta}\right) = \left(1 - \frac{\Delta_1}{\Delta}\right)\left(u'\left(1 - \frac{\Delta_1}{\Delta}\right) - \underline{\theta}\right) + \frac{\nu}{1-\nu}\Delta_1.$$

Clearly, the only possibly profitable deviation from this candidate equilibrium consists in offering the pooling contract $w_1^D = \overline{\theta}$ in the first period,[18] and the separating

[18] To see this, notice that, since the agents' mixed strategy γ^* does not depend on the wages w_1 and w_2, but only on the equilibrium price $p_2 = u'(\nu)$, it follows that, for every value of the wages, agents

contract $w_2^D = \underline{\theta}$ in the second period.[19]. Accordingly, the deviation profit is

$$\pi^D \triangleq u'(\nu\gamma^*) - \overline{\theta} + \nu(u'(\nu) - \underline{\theta}) = u'\left(1 - \frac{\Delta_1}{\Delta}\right) - \underline{\theta} - \Delta + \frac{\nu}{1-\nu}\Delta_1.$$

It can be immediately checked that $\pi^{SS} = \pi^D$ if and only if Δ solves equation (3.3). Since the left-hand side and the right-hand side of equation (3.3) are, respectively, increasing and decreasing in Δ, it follows that the symmetric semi-separating equilibrium exists for all $\Delta \geq \Delta_3$, whenever Δ_3 satisfies Assumption 3.1 — i.e., if $\Delta_3 \leq \frac{\Delta_0}{1-\nu}$. ■

Proof of Proposition 3.5. Clearly, aggregate output cannot increase over time when all principals offer the pooling contract in the first period — i.e., for all $\Delta < \Delta_2$. On the contrary, for all $\Delta > \Delta_3$: $y_1^* = \nu\gamma^* < y_2^* = \nu$.

Next, it is easy to obtain that, for all $\Delta \in (\Delta_2, \Delta_3)$,

$$y_1^* = u'^{-1}\left(\underline{\theta} + \frac{\Delta^2}{\Delta_1}\right), \qquad y_2^* = \nu. \tag{3.8}$$

Moreover, y_1^* is clearly decreasing in Δ. Finally, $y_1^*|_{\Delta\to\Delta_2} = 1 > y_2^*$ and $y_1^*|_{\Delta\to\Delta_3} = \nu\gamma^*|_{\Delta\to\Delta_3} = 1 - \frac{\Delta_1}{\Delta} < y_2^*$, which concludes the proof. ■

Proof of Proposition 3.6. It can be easily checked that, for $\Delta \in (\Delta_2, \Delta_3]$, the equilibrium profit is

$$\pi^* \triangleq \frac{\Delta^2}{\Delta_1} - \Delta + \frac{\nu}{1-\nu}\Delta_1.$$

Hence, $\frac{\partial \pi^*}{\partial \Delta} > 0 \iff \Delta > \frac{\Delta_1}{2}$, which is satisfied in the considered range of the parameters.

Consumer welfare, in our simple competitive environment (and no discounting)[20], is defined as

$$W \triangleq u(y_1) - y_1 u'(y_1) + u(y_2) - y_2 u'(y_2),$$

and it is an increasing function of aggregate outputs.[21] Thus, since, from (3.8), y_1^* is

do not have any profitable deviation from their mixed strategy. Moreover, since we proved that each principal's optimal choice within the set $w_1 \in [\underline{\theta}, \overline{\theta})$ is $w_1^* = \underline{\theta}$, a principal could profitably deviate from the equilibrium strategy in the first period only by offering the pooling contract $w_1^D = \overline{\theta}$.

[19] Given that $p_2 = u'(\nu)$, and $\Delta > \Delta_1$, the deviating principal (who did not learn the agent's type in the first period) finds it optimal to offer $w_2^D = \underline{\theta}$.

[20] Clearly, since, in the equilibrium under consideration, the second period output does not depend on Δ, our comparative statics result holds true even assuming that consumers, differently from players, discount future utility.

[21] For all $t = 1, 2 : \frac{\partial W}{\partial y_t} = -y_t u''(y_t) > 0$, by the concavity of the utility function.

decreasing in Δ, whereas y_2^* does not depend on Δ, it follows $\frac{\partial W}{\partial \Delta} < 0$. ■

Proof of Proposition 3.7. The proof proceeds as follows: through a series of Lemmata, we establish that symmetric pure strategy equilibria and mixed strategy equilibria in which agents randomize do not exist for $\Delta \in (\Delta_0, \Delta_1)$.

Lemma 3.A.1. *A (symmetric) pure strategy equilibrium in which all principals offer the pooling contract in both periods (i.e., $w_1^* = w_2^* = \overline{\theta}$), exists if and only if $\Delta \leq \Delta_0$. Moreover, there are no other equilibria in which all principals offer the pooling contract in the first period.*

Proof. Consider a symmetric equilibrium in which all principals offer the pooling contract in the first period ($w_1^* = \overline{\theta}$). Clearly, in such candidate equilibrium, each principal finds it optimal to offer the pooling contract also in the second period whenever $\Delta \leq \Delta_0$. In this case, her expected profit would be

$$\pi^{PP} \triangleq (1+\delta)(u'(1) - \overline{\theta}).$$

Consider the deviation consisting in offering the separating contract in the first period ($w_1^D = \underline{\theta} + \delta\Delta$). Since the deviating principal then achieves full rent extraction in the second period, she obtains

$$\pi^D \triangleq \nu(u'(1) - \overline{\theta}) + \delta[u'(1) - \nu\underline{\theta} - (1-\nu)\overline{\theta}] = \nu(u'(1) - \underline{\theta}) + \delta(u'(1) - \overline{\theta}),$$

where $\pi^P \geq \pi^D$ if and only if $\Delta \leq \Delta_0$. Notice that there are no other possible deviations from this candidate equilibrium, since, as we already shown in the case with no discounting, the semi-separating outcome cannot be obtained for $\Delta \leq \Delta_0$.

For $\Delta_0 < \Delta < \Delta_1$, in a candidate equilibrium in which each principal offers the pooling contract in the first period, the mixed strategy equilibrium defined by (3.1) is played in the second period. Accordingly, each principal's profit is

$$\pi^{PM} \triangleq \underbrace{u'(1) - \overline{\theta}}_{w_1 = \overline{\theta}} + \underbrace{+\delta\left[\alpha^*\nu\left(u'\left(\nu\alpha^* + 1 - \alpha^*\right) - \underline{\theta}\right) + (1-\alpha^*)\left(u'\left(\nu\alpha^* + 1 - \alpha^*\right) - \overline{\theta}\right)\right]}_{\mathbb{E}[w_2] = \alpha^*\underline{\theta} + (1-\alpha^*)\overline{\theta}}.$$

A deviation for any principal consists in offering the separating contract in the first period, thus being able to achieve the first-best outcome in the second period, and obtaining

$$\pi^D \triangleq \nu((u'(1) - \underline{\theta} - \delta\Delta) + \delta(u'(\nu\alpha^* + 1 - \alpha^*) - \underline{\theta})) + (1-\nu)\delta\underbrace{(u'(\nu\alpha^* + 1 - \alpha^*) - \overline{\theta})}_{=\nu(u'(\nu\alpha^* + 1 - \alpha^*) - \underline{\theta})},$$

It can be easily checked that $\pi^{PM} \geq \pi^{D}$ if and only if

$$\Delta_0 - (1 - \delta\alpha^*)\Delta - \delta\alpha^*(1-\nu)(u'(\alpha^*\nu + 1 - \alpha^*) - \underline{\theta}) > 0,$$

which, substituting α^* from (3.1), is satisfied for all $\Delta < \Delta_0$: thus, the considered equilibrium does not exist.

Lastly, consider $\Delta \geq \Delta_1$: in this case, in a symmetric equilibrium in which all principals offered the pooling contract in the first period, they find it optimal to offer the separating contract in the second period. Accordingly, each principal's profit is given by

$$\pi^{PS} \triangleq u'(1) - \overline{\theta} + \delta\nu(u'(\nu) - \underline{\theta}).$$

Again, a deviation for any principal consists in offering the separating contract in the first period, thus being able to achieve the first-best outcome in the second period, and obtaining

$$\pi^{D} \triangleq \nu[u'(1) - \underline{\theta} - \delta\Delta] + \delta[u'(\nu) - \nu\underline{\theta} - (1-\nu)\overline{\theta}] = \nu(u'(1) - \underline{\theta}) + \delta(u'(\nu) - \overline{\theta}).$$

Hence, $\pi^{P} \geq \pi^{D}$ if and only if

$$\Delta \leq \frac{\Delta_0 - \delta\Delta_1}{1-\delta} < \Delta_1.$$

Thus, the considered equilibrium does not exist. ∎

Lemma 3.A.2. *A necessary condition for the existence of a (symmetric) pure strategy equilibrium in which all principals offer the separating contract $w_1^* = \underline{\theta} + \delta\Delta$ in the first period (hence, full surplus extraction is obtained in the second period) is $\Delta \geq \frac{\Delta_1 - \delta\Delta_0}{1-\delta} > \Delta_1$.*

Proof. Consider a symmetric equilibrium in which only efficient types produce in the first period — i.e., each principal offers the separating contract in $t = 1$. In such an equilibrium, $w_1^* = \underline{\theta} + \delta\Delta < \overline{\theta}$, $w_2^* = \overline{\theta}$ if an agent has rejected w_1^* in the first period, and $w_2^* = \underline{\theta}$ otherwise. It is immediate to see that an efficient type has no incentive to reject the first period offer. Hence, there is full surplus extraction in the second period and every principal obtains an equilibrium profit

$$\pi^{S} \triangleq \nu\left[(u'(\nu) - \underline{\theta} - \delta\Delta) + \delta(u'(1) - \underline{\theta})\right] + (1-\nu)\,\delta\left(u'(1) - \overline{\theta}\right).$$

A first possible deviation from this candidate equilibrium for any principal consists in offering the pooling contract $w_1^D = \overline{\theta}$ in the first period. In the second period, for $\Delta < \Delta_0$, the deviating principal finds it optimal to offer the pooling contract, getting

$$\pi^D \triangleq u'(\nu) - \underline{\theta} - \Delta + \delta\left(u'(1) - \overline{\theta}\right).$$

Comparing the equilibrium and the deviation profits, we find

$$\pi^S \geq \pi^D \iff \Delta \geq \Delta_1.$$

Thus, for $\Delta < \Delta_0$, we have found a profitable deviation: the considered strategy profile does not constitute an equilibrium.

Analogously, for $\Delta \geq \Delta_0$, the profit of a deviating principal who offers the pooling contract in the first period and the separating contract in the second period is

$$\pi^D \triangleq u'(\nu) - \overline{\theta} + \nu\delta\left[u'(1) - \underline{\theta}\right].$$

Comparing the equilibrium and the deviation profits, we find

$$\pi^S \geq \pi^D \iff \Delta \geq \frac{\Delta_1 - \delta\Delta_0}{1-\delta} > \Delta_1,$$

which thus constitutes a necessary condition for the existence of this equilibrium.[22] ■

Lemma 3.A.3. *In a (symmetric) semi-separating equilibrium, each principal offers $w_1^* = w_2^* = \underline{\theta}$. An efficient agent accepts w_1^* with probability*

$$\gamma^* = \frac{\Delta - \Delta_1}{\nu\Delta}$$

and always produces in the second period. Inefficient types are shut down in both periods. A necessary condition for the existence of this equilibrium is $\Delta > \Delta_1$.

Proof. In the (symmetric) semi-separating equilibrium, each principal offers $w_2^* = \underline{\theta}$, to an agent who has not accepted the wage $w_1 \in [\underline{\theta}, \underline{\theta} + \delta\Delta)$, with a probability σ^* such that the efficient agent is indifferent between accepting or not the wage w_1 — i.e.,

$$w_1 - \underline{\theta} = \delta(1-\sigma^*)\Delta \iff \sigma^* = 1 - \frac{w_1 - \underline{\theta}}{\delta\Delta}.$$

As in the case without discounting, each efficient agent accepts w_1 with a probability γ^* such that the principal is indifferent between offering $w_2 = \underline{\theta}$ and $w_2 = \overline{\theta}$ to an agent who

[22] Such a condition is not in general sufficient for the existence of the considered equilibrium, since for all $\Delta > \Delta_1$ a deviating principal could also offer $w_1 = \underline{\theta}$ so to induce the semi-separating outcome.

refused w_1 — i.e., for any given market price p_2, γ^* is given by (3.7), which is positive as long as

$$\Delta > (1-\nu)(p_2 - \underline{\theta}). \tag{3.9}$$

Such inequality is thus a necessary condition for the existence of the symmetric semi-separating equilibrium.

Moreover, by the same argument outlined in the case with $\delta = 1$, it can be easily obtained that $w_1^* = \underline{\theta}$ and, as a consequence, $\sigma^* = 1$, hence (with probability one) $w_2^* = \underline{\theta}$. Thus, only the efficient agents produce in the second period, yielding $p_2 = u'(\nu)$. Therefore, the necessary condition (3.9) for the existence of this equilibrium becomes $\Delta > \Delta_1$. ■

Lemma 3.A.4. *A necessary condition for the existence of any (symmetric) mixed strategy equilibrium in which the semi-separating contract is offered with positive probability is* $\Delta > \Delta_1$.

Proof. In what follows, we derive the existence conditions for every mixed strategy equilibrium in which principals offer the semi-separating contract with positive probability.

Randomization between pooling and semi-separating contracts. Consider a mixed strategy symmetric equilibrium, in which, in $t = 1$, principals randomize between the semi-separating contract ($w_1 = \underline{\theta}$) and the pooling contract ($w_1 = \underline{\theta} + \Delta$). Then, at time $t = 2$, all principals offer $w_2 = \underline{\theta}$. This is obviously the case for the principals who offered $w_1 = \underline{\theta}$. To see why this holds also for the principals who offered the pooling contract, notice that the semi-separating outcome can be achieved if (3.9) holds, and under the same condition the principals who offered the pooling contract in $t = 1$ find it optimal to offer the separating contract in $t = 2$. Given that $p_2^* = u'(\nu)$, it follows from (3.9) that this candidate equilibrium can be defined only for $\Delta > \Delta_1$. ■

Randomization between separating and semi-separating contracts. Consider a mixed strategy symmetric equilibrium, in which, in $t = 1$, principals randomize between the semi-separating contract ($w_1 = \underline{\theta}$) and the separating contract ($w_1 = \underline{\theta} + \delta\Delta$). Say that $w_1 = \underline{\theta}$ is offered with probability υ. Clearly, at time $t = 2$, the principals who offered $w_1 = \underline{\theta}$ offer $w_2 = \underline{\theta}$, whereas the principals who offered the separating contract achieve the first-best outcome. It can be easily obtained that principals are indifferent between offering the separating and the semi-separating contract at time $t = 1$ if

$$(1-\nu)(u'(\upsilon\nu\gamma^* + (1-\upsilon)\nu) - \underline{\theta}) \left(\frac{u'(\upsilon\nu + 1 - \upsilon) - \underline{\theta}}{\Delta} - 1 \right) + \delta(1-\nu)(u'(\upsilon\nu + 1 - \upsilon) - \underline{\theta}) - \delta\Delta = 0,$$

whose left-hand side is increasing in υ: hence, in order for the considered candidate equilibrium to exist, the left-hand side of the above equation, when evaluated at $\upsilon = 0$, must be negative. Such condition leads to

$$\Delta > \frac{\delta\Delta_0 - \Delta_1}{2\delta} + \frac{\sqrt{(1-\nu)(\Delta_1 - \delta\Delta_0)^2 + 4\delta\Delta_0\Delta_1}}{2\delta\sqrt{1-\nu}},$$

this threshold being higher than Δ_1 if

$$\frac{\Delta_1}{\Delta_0} < \frac{\delta(1-\nu)+1}{(1-\nu)(1+\delta)},$$

which is satisfied under Assumption 3.2. $\blacksquare$

Randomization among pooling, separating and semi-separating contracts. Finally, consider a symmetric equilibrium in which, in $t = 1$, principals offer the semi-separating contract ($w_1 = \underline{\theta}$) with probability μ, the separating contract ($w_1 = \underline{\theta} + \delta\Delta$) with probability κ, and the pooling contract ($w_1 = \underline{\theta} + \Delta$) otherwise (i.e., with probability $1 - \mu - \kappa$). Then, at time $t = 2$, the principals who offered the semi-separating contract offer $w_2 = \underline{\theta}$, the principals who offered the separating contract obtain the first-best outcome, whereas the principals who offered the pooling contract offer the separating contract ($w_2 = \underline{\theta}$).
After some algebra, it can be obtained that principals are indifferent among these three alternative strategy profiles if and only if the following system is satisfied

$$\begin{cases} \Delta^2 = (1-\nu)(p_1^* - \underline{\theta})(p_2^* - \underline{\theta}) \\ \Delta = \frac{1-\nu}{1-\delta}(p_1^* - \underline{\theta} - \delta(p_2^* - \underline{\theta})) \end{cases}$$

where equilibrium prices are

$$p_1^* = u'(\mu\nu\gamma^* + \kappa\nu + 1 - \mu - \kappa), \quad p_2^* = u'(\nu + \kappa(1-\nu)),$$

with γ^* being defined by (3.7).
From the second equation we find $p_1^* - \underline{\theta} = \frac{1-\delta}{1-\nu}\Delta + \delta(p_2^* - \underline{\theta})$, which we substitute into the first equation to find

$$\Delta^2 = ((1-\nu)\delta(p_2^* - \underline{\theta}) + (1-\delta)\Delta)(p_2^* - \underline{\theta}),$$

which can be solved for p_2^* yielding the only positive solution:

$$p_2^* = \underline{\theta} + \Delta\frac{\sqrt{(1+\delta)^2 - 4\delta\nu} - 1 + \delta}{2\delta(1-\nu)},$$

from which we easily obtain

$$\kappa^* = \frac{1}{1-\nu}\left(u'^{-1}\left(\overline{\theta} + \Delta\frac{\sqrt{(1+\delta)^2 - 4\delta\nu} - 1 + \delta}{2\delta(1-\nu)}\right) - \nu\right).$$

Hence

$$\kappa^* \in (0,1) \iff \Delta \in \left(\frac{2\delta\Delta_0}{\sqrt{(1+\delta)^2 - 4\delta\nu} - 1 + \delta}, \frac{2\delta\Delta_1}{\sqrt{(1+\delta)^2 - 4\delta\nu} - 1 + \delta}\right),$$

and the lower bound of the interval is higher than Δ_1 under Assumption 3.2. ■

Proof of Proposition 3.8. To establish the result, we have to characterize the symmetric mixed strategy equilibria in which principals randomize between offering the separating and the pooling contract, given that, from Proposition 3.7, we already know that no other kind of equilibria can exist in the considered region of parameters.
Thus, consider symmetric equilibria in which, in $t = 1$, each principal offers the separating contract with probability β, and the pooling contract otherwise. We thus have $p_1 = u'(\nu\beta + 1 - \beta)$. Clearly, the principals who offered the separating contract in $t = 1$, then achieve the first-best outcome in $t = 2$.
As for the principals who proposed the pooling contract in $t = 1$, first notice that they never find it optimal to offer the pooling contract in the second period. In fact, if they offer the pooling contract also in the second period, then the market price will be $p_2 = u'(1)$, which immediately implies that this choice is never optimal for all $\Delta > \Delta_0$.

Next consider the case in which the principals who offered the pooling contract in $t = 1$ offer the separating contract in the second period. In this case, the market price in $t = 2$ is $p_2 = u'(\beta + (1-\beta)\nu)$. Consequently, this strategy constitutes an optimal choice for these principals if $\Delta \geq (1-\nu)(u'(\beta + (1-\beta)\nu) - \underline{\theta})$. In this case, a principal's expected profit if she offers the separating contract in the first period (hence, the first-best solution is achieved in the second period) is

$$\pi^S \triangleq \nu(u'(\beta\nu + 1 - \beta) - \underline{\theta} - \delta\Delta) + \delta(u'(\beta + (1-\beta)\nu) - \underline{\theta} - (1-\nu)\Delta),$$

whereas, if she offers the pooling contract in the first period (and the separating contract in the second period), she obtains

$$\pi^{PS} \triangleq u'(\beta\nu + 1 - \beta) - \underline{\theta} - \Delta + \delta\nu(u'(\beta + (1-\beta)\nu) - \underline{\theta}).$$

It can be easily seen that $\pi^S = \pi^{PS}$ if and only if equation (3.6) is satisfied. Moreover, the right-hand side of equation (3.6) is increasing in β^*, from which it follows that such mixed strategy should be an increasing function of Δ. Next notice that, since we are considering $\Delta \geq (1-\nu)(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta})$, this equilibrium can exists if

$$\frac{1-\nu}{1-\delta}(u'(\beta^*\nu + 1 - \beta^*) - \underline{\theta} - \delta(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta})) > (1-\nu)(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta}),$$

which gives $\beta^* > \frac{1}{2}$.
Hence, there exists a (unique) equilibrium value β^* satisfying (3.6) if

$$\Delta - \left(\frac{1-\nu}{1-\delta}(u'(\beta\nu + 1 - \beta) - \underline{\theta} - \delta(u'(\beta + (1-\beta)\nu) - \underline{\theta}))\right)\Bigg|_{\beta=1} < 0$$

and

$$0 < \Delta - \left(\frac{1-\nu}{1-\delta}(u'(\beta\nu + 1 - \beta) - \underline{\theta} - \delta(u'(\beta + (1-\beta)\nu) - \underline{\theta}))\right)\Bigg|_{\beta=\frac{1}{2}},$$

yielding

$$\Delta \in \left(\Delta^*, \Delta_0 + \frac{\Delta_1 - \Delta_0}{1-\delta}\right).$$

It is straightforward to see that $\Delta^* \in (\Delta_0, \Delta_1)$, whereas $\Delta_0 + \frac{\Delta_1 - \Delta_0}{1-\delta} > \Delta_1$. However, to establish the existence of the considered candidate equilibrium, we still need to check whether there are profitable deviations from it. Clearly, the only possibly profitable deviation for each principal consists in offering the semi-separating contract.[23] The deviating principal would obtain

$$\pi^D \triangleq \nu\gamma^*(u'(\nu\beta^* + 1 - \beta^*) - \underline{\theta}) + \delta\nu(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta}),$$

with γ^* being given by (3.7). After simple algebra, it follows that the equilibrium profit $\pi^{PS}(= \pi^S)$ is higher than π^D if and only if

$$\Delta^2 < (1-\nu)(u'(\nu\beta^* + 1 - \beta^*) - \underline{\theta})(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta}). \tag{3.10}$$

Since the right-hand side is decreasing[24] in Δ, it follows that the above inequality is

[23]From (3.9), we know that the semi-separating outcome is achievable if $\Delta \geq (1-\nu)(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta})$, which is always the case in the considered region of parameters. To see this, notice that the right-hand side of this inequality is decreasing in Δ, hence a sufficient condition for this inequality to hold for all Δ in the considered interval is that it is satisfied for $\Delta \to \Delta^*$, which can be immediately verified.

[24]The derivative of the right-hand side with respect to Δ is negative if

$$u''(\nu\beta^* + 1 - \beta^*)(u'(\beta^* + (1-\beta^*)\nu) - \underline{\theta}) > u''(\beta^* + (1-\beta^*)\nu)(u'(\nu\beta^* + 1 - \beta^*) - \underline{\theta}).$$

Since $u'(\beta^* + (1-\beta^*)\nu) \leq u'(\nu\beta^* + 1 - \beta^*)$, a sufficient condition in order for the above inequality to hold is that $u''(\nu\beta^* + 1 - \beta^*) \leq u''(\beta^* + (1-\beta^*)\nu)(< 0)$ — i.e., that $u''(\cdot)$ is a decreasing function, hence satisfied for Δ sufficiently small, say $\Delta < \overline{\Delta}$. To prove that $\overline{\Delta} > \Delta_1$, notice that the right-hand side of inequality (3.10) is higher than $\frac{\Delta_0\Delta^*}{1-\nu}$, thus, a sufficient condition for (3.10) to hold at $\Delta = \Delta_1$ is $\Delta_0 > (1-\nu)\frac{\Delta_1^2}{\Delta^*}$, which is the case under Assumption 3.2.[25]

Finally, we have to consider the case in which the principals who in the first period offered the pooling contract, then randomize again in the second period between offering the separating and the pooling contract. Denote by ϵ the probability of offering $w_2 = \underline{\theta}$. Then, the market price in $t = 2$ is $p_2 = u'(\beta + (1-\beta)(\epsilon\nu + 1 - \epsilon))$. Thus, a principal who offered the pooling contract in the first period is indifferent between offering the pooling and the separating contract in $t = 2$ if

$$\nu(u'(\beta + (1-\beta)(\epsilon\nu + 1 - \epsilon)) - \underline{\theta}) = u'(\beta + (1-\beta)(\epsilon\nu + 1 - \epsilon)) - \underline{\theta} - \Delta,$$

which yields

$$\varepsilon^*(\beta) = \frac{1}{1-\nu}\left(1 - \frac{u'^{-1}\left(\frac{\Delta}{1-\nu} + \underline{\theta}\right) - \beta}{1-\beta}\right). \tag{3.11}$$

We have $\frac{\partial\varepsilon^*}{\partial\beta} > 0$ for all $\Delta > \Delta_0$, hence $\epsilon^*(\beta)$ takes its lower value at $\beta = 0$: $\epsilon^*|_{\beta=0} = \frac{1}{1-\nu}\left[1 - u'^{-1}\left(\frac{\Delta}{1-\nu} + \underline{\theta}\right)\right] > 0 \iff \Delta > \Delta_0$, and $\epsilon^*(\beta) \leq 1$ for

$$\beta \leq \bar{\beta} \triangleq \frac{1}{1-\nu}\left[u'^{-1}\left(\frac{\Delta}{1-\nu} + \underline{\theta}\right) - \nu\right].$$

Notice that, regardless of the equilibrium value of β, the market price in the second period is $p_2 = \frac{\Delta}{1-\nu} + \underline{\theta}$. Thus, if a principal offers the pooling contract in the first period, and plays according to the mixed strategy defined by ϵ^* in the second period, her expected profit is

$$\pi^{PM} \triangleq u'(\beta\nu + 1 - \beta) - \underline{\theta} - \Delta + \delta\left(\epsilon^*\nu\frac{\Delta}{1-\nu} + (1-\epsilon^*)\left(\frac{\Delta}{1-\nu} - \Delta\right)\right),$$

whereas, if she offers the separating contract in the first period (hence, the first-best outcome is achieved in the second period), she obtains

$$\pi^S \triangleq \nu(u'(\beta\nu + 1 - \beta) - \underline{\theta} - \delta\Delta) + \delta\left(\frac{\Delta}{1-\nu} - (1-\nu)\Delta\right).$$

By equating the two above payoffs, we find that the equilibrium value is given by $\beta = \alpha^*$,

---

$u'''(\cdot) \leq 0$, as we assumed in Section 3.5.

[25]Specifically, the considered inequality is implied by Assumption 3.2 if $(1-\nu)\frac{\Delta_1}{\Delta^*} < \sqrt{1-\nu}$, which is in turn implied by Assumption 3.2. In fact, the considered inequality can be rewritten as $\frac{u'(\nu)-\underline{\theta}}{u'\left(\frac{1+\nu}{2}\right)-\underline{\theta}} < \frac{1}{\sqrt{1-\nu}}$, whose left-hand side is smaller than $\frac{\Delta_1}{\Delta_0}$.

as defined by (3.1), and it can be easily checked that $\alpha^* \in (0, \bar{\beta}) \iff \Delta \in (\Delta_0, \Delta^*)$. Finally, by substituting the equilibrium value α^* into (3.11), we get the equilibrium value $\epsilon^*(\alpha^*)$ given by (3.5). Notice that there are not possibly profitable deviations from this equilibrium, since it cannot be optimal to offer $w_1^D = \underline{\theta}$, given that no agent would accept the first period offer.[26] ■

Proof of Proposition 3.9. It is easy to find that, for $\Delta \in (\Delta_0, \Delta^*)$: $y_1^* = y_2^* = u'^{-1}\left(\underline{\theta} + \frac{\Delta}{1-\nu}\right)$, whereas, for $\Delta \in (\Delta^*, \Delta_1)$: $y_1^* = \beta^*\nu + 1 - \beta^* < y_2^* = \beta^* + (1-\beta^*)\nu$ (since $\beta^* > \frac{1}{2}$). ■

Proof of Proposition 3.10. The result concerning the first-period quantities is trivial, since $y_1^*|_{\delta=1} = 1$, whereas not every principal produces when $\delta < 1$. As for the second period, we already noticed that, when $\Delta < \Delta^*$, in both games the aggregate output coincides with the equilibrium output of the static game. Finally, for $\Delta > \Delta^*$, we have:

$$y_2^*|_{\delta=1} < y_2^*|_{\delta<1} \iff \alpha^* + \beta^* > 1,$$

which is satisfied if $\alpha^* > \beta^*$ (recall that $\beta^* > \frac{1}{2}$). To see that this is always the case, notice that

$$\frac{1-\delta}{1-\nu}\Delta < u'(\nu\alpha^* + 1 - \alpha^*) - \underline{\theta} - \delta(u'(\alpha^* + (1-\alpha^*)\nu) - \underline{\theta}) \iff \Delta > \Delta^*.$$

That is, the right-hand side of equation (3.6), which is increasing in β, when evaluated at $\beta = \alpha^*$ is higher than the left-hand side, which proves that $\alpha^* > \beta^*$. ■

Proof of Proposition 3.11. In the mixed strategy equilibrium defined for $\Delta \in (\Delta_0, \Delta^*]$, principals' equilibrium profit is

$$\pi_1^* \triangleq \frac{\nu(1+\delta)}{1-\nu}\Delta,$$

which, interestingly, does not depend on the demand function, and is clearly increasing in Δ. In the mixed strategy equilibrium defined for $\Delta \in (\Delta^*, \Delta_1]$, the equilibrium profit can be written as

$$\pi_2^* \triangleq \nu(u'(1-(1-\nu)\beta^*) - \underline{\theta} - \delta\Delta) + \delta(u'(\nu + (1-\nu)\beta^*) - \underline{\theta} - (1-\nu)\Delta),$$

hence

$$\frac{\partial \pi_2^*}{\partial \Delta} = (1-\nu)\frac{\partial \beta^*}{\partial \Delta}\left(\delta u''(\nu + (1-\nu)\beta^*) - \nu u''(1-(1-\nu)\beta^*)\right) - \delta,$$

[26] In fact, substituting the equilibrium price into (3.7), it can be easily seen that $\gamma^* = 0$, which implies that the semi-separating outcome cannot be achieved by a deviating principal.

where, as proved above, $\frac{\partial \beta^*}{\partial \Delta} \geq 0$. It can be immediately seen that

$$\left.\frac{\partial \pi_2^*}{\partial \Delta}\right|_{\delta \to 0} = -\nu(1-\nu)\left.\frac{\partial \beta^*}{\partial \Delta}\right|_{\delta \to 0} u''(1-(1-\nu)\beta^*) > 0.$$

Moreover, from (3.6), it can be seen that, as $\delta \to 1$: $\beta^* \to \frac{1}{2}$, implying that

$$\left.\frac{\partial \pi_2^*}{\partial \Delta}\right|_{\delta \to 1} = (1-\nu)^2\left.\frac{\partial \beta^*}{\partial \Delta}\right|_{\delta \to 1} u''\left(\frac{1+\nu}{2}\right) - 1 < 0.$$

We finally consider consumer welfare, defined as[27]

$$W \triangleq u(y_1) - y_1 u'(y_1) + \delta(u(y_2) - y_2 u'(y_2)).$$

Differentiating it with respect to Δ yields

$$\frac{\partial W}{\partial \Delta} = -y_1 \frac{\partial y_1}{\partial \Delta} u''(y_1) - \delta y_2 \frac{\partial y_2}{\partial \Delta} u''(y_1).$$

For $\Delta \in (\Delta_0, \Delta^*]$, since $y_1^* = y_2^*$ is decreasing in Δ, we can immediately conclude $\frac{\partial W^*}{\partial \Delta} < 0$. In the equilibrium defined for $\Delta \in (\Delta^*, \Delta_1]$,

$$\frac{\partial y_1^*}{\partial \Delta} = -(1-\nu)\frac{\partial \beta^*}{\partial \Delta} = -\frac{\partial y_2^*}{\partial \Delta},$$

from which we can write

$$\frac{\partial W^*}{\partial \Delta} = y_1^* u''(y_1^*) - \delta y_2^* u''(y_1^*).$$

Thus, a sufficient condition in order for $\frac{\partial W^*}{\partial \Delta} < 0$ is $y_1^* u''(y_1^*) < y_2^* u''(y_1^*)$, which is in turn always satisfied when $u'''(\cdot) < 0$. ■

Appendix 3.B. Long-term contracts and renegotiation

In this Appendix, we study how the equilibrium configuration of the game changes when we introduce the possibility of commitment and Pareto-improving renegotiation — i.e., when principals can commit to a long-term contract, offered at the beginning of the game and which rules the relationship in both periods, however they cannot commit not to offer a new contract at the beginning of the second period that would replace the initial contract if the agent finds it acceptable. We first consider the model without discounting.

[27] For coherence, we have considered the discount factor δ (used by principals and agent) also for consumers. However, it can be immediately verified that this assumption is immaterial to our results.

In our model, long-term contracts are as follows:[28]

(1) contracts specifying that $w_1 = \underline{\theta}$ and

 (a) $w_2 = \overline{\theta}$ for every $x_1 \in \{0, 1\}$,

 (b) $w_2 = \underline{\theta}$ for every $x_1 \in \{0, 1\}$,

 (c) $w_2 = \underline{\theta}$ if $x_1 = 1$ and $w_2 = \overline{\theta}$ otherwise,

 (d) $w_2 = \overline{\theta}$ if $x_1 = 1$ and $w_2 = \underline{\theta}$ otherwise;

(2) contracts specifying that $w_1 = \overline{\theta}$ and w_2 is as in the cases (a),(b),(c) and (d) above.

Notice that contracts in which $w_2 = \overline{\theta}$ for every $x_1 \in \{0, 1\}$, namely those labeled with (1a) and (2a), are clearly renegotiation-proof. Importantly, contract (1a) implements the separating outcome of the model with spot contracts, since in the first period only efficient agents produce whereas in the second period all agents produce, and the inter-temporal rent of efficient types is Δ. Thus, differently from a model with spot contracting, long-term contracting gets rid of the ratchet effect, since the principal can commit to offer $w_2 = \overline{\theta}$ to an agent who produced when paid $w_1 = \underline{\theta}$. By the same token, contract (2a) clearly implements the pooling outcome of the model with spot contracts, since in both periods all agents produce and efficient types obtain a rent Δ in either period. On the contrary, contracts in which $w_2 = \underline{\theta}$, for some x_1, are not renegotiation proof if a principal, after observing the first-period outcome, would be better off by offering $w_2 = \overline{\theta}$ — i.e., if $\Delta < (1 - \Pr[\theta = \underline{\theta} | x_1 = 0])(p_2 - \underline{\theta})$. Thus, contract (1b) is renegotiation-proof only when efficient agents randomize — i.e., when the semi-separating outcome is implemented.[29,30] Contract (2b) is renegotiation proof if $\Delta \geq (1 - \nu)(p_2 - \underline{\theta})$. The other contracts can be neglected: contract (1c) is not incentive-compatible;[31] contract (1d) is not renegotiation proof;[32] contract (2c) is equivalent to contract (2b) or to contract (1a);[33] contract (2d) is

[28] Clearly, each period wage is paid conditional on the agent producing in that period.

[29] To see this, notice that, if all efficient agents accept to produce in the first period being paid $w_1 = \underline{\theta}$, then $\Pr[\theta = \underline{\theta} | x_1 = 0] = 0$ and, by Assumption 3.1 contract (1b) is not renegotiation proof. Obviously, if all efficient agents decide not to produce in the first period, we can neglect contract (1b) since principals never find it optimal to offer it.

[30] To be more explicit, contract (1b) can be seen as a contract in which the agent has three options in total: in the first period, he can choose between producing or not; if he produces in the first period, he is to produce in the second period also, and obtains a wage $\underline{\theta}$ each period; if he does not produce in the first period, he then has a second choice in the second period: either he produces, and obtains a wage $\underline{\theta}$, or he does not produce.

[31] In fact, efficient agents would be better off not producing in the first period, which clearly harms principals.

[32] If an agent did not produced in the first period, at the beginning of the second period the principal knows that he is inefficient, hence offering $w_2 = \overline{\theta}$ would be Pareto-improving.

[33] Indeed, under contract (2c), efficient types always accept w_1, whereas inefficient types are indifferent

equivalent to contract (2a).[34]

To sum up, in the game without discounting, for every outcome (i.e., market prices and inter-temporal rents and profits) that can be achieved by means of short-term contracts there exists a long-term renegotiation-proof contract which yields the same outcome. However, there exists a long-term renegotiation proof contract which allows principals to achieve full separation in the first period, which is an outcome that, in the game with no discounting, cannot be achieved by using spot contracts. As a consequence, the equilibrium set of the game under long-term renegotiation proof contracts differs from the one characterized in Section 3.4. Specifically, we are able to prove the following result.

Proposition 3.B.1. *The equilibria of the game with long-term renegotiation-proof contracts are as follows.*

- *If $\Delta \leq \Delta_0$, in the unique symmetric equilibrium of the game, each principal offers contract (2a).*
- *If $\Delta \in (\Delta_0, \Delta^*]$, there exists only one symmetric equilibrium in which, with probability α^*, defined in (3.1), each principal offers contract (1a), with the same probability α^*, each principal offers contract (2b), and with complementary probability $1 - 2\alpha^*$, each principal offers contract (2a).*
- *If $\Delta \in \left(\Delta^*, \frac{\Delta^*}{\sqrt{1-\nu}}\right)$ there exists only one symmetric equilibrium in which principals offer with the same probability ($\frac{1}{2}$) contracts (1a) and (2b).*
- *If $\Delta \in \left(\frac{\Delta^*}{\sqrt{1-\nu}}, \hat{\Delta}\right]$, where*
$$\hat{\Delta} \triangleq \frac{\Delta_1}{\sqrt{1-\nu}},$$
there exists only one symmetric equilibrium in which principals offer contract (1a) with probability
$$\phi^* \triangleq \frac{1}{1-\nu}\left(u'^{-1}\left(\overline{\theta} + \frac{\Delta}{\sqrt{1-\nu}}\right) - \nu\right),$$
contract (1b) with probability
$$\varphi^* \triangleq \frac{1}{\sqrt{1-\nu}}\left(1 + \nu - 2u'^{-1}\left(\overline{\theta} + \frac{\Delta}{\sqrt{1-\nu}}\right)\right)$$
and contract (2b) otherwise.

---

between accepting the first period offer (thus do not producing in the second period) or declining the first period offer and produce in the second period. In the former case, this contract is equivalent to contract (2b), in the latter case is equivalent to contract (1a).

[34] In fact, all agents produce in the first period, given that $w_2 = \overline{\theta}$ if $x_1 = 1$.

- *If $\Delta \in \left(\hat{\Delta}, \Delta_3\right]$, there exists only one symmetric equilibrium in which principals offer contract (1b) with probability ρ^* defined by (3.4), and contract (2b) otherwise.*

- *If $\Delta > \Delta_3$, in the unique symmetric equilibrium of the game, each principal offers contract (1b).*

Proof. To begin with, we calculate the profit that, for any market prices p_1 and p_2, a principal obtains offering each one of the four relevant long-term contracts:

$$\begin{aligned}
\pi^{(1a)} &\triangleq \nu(p_1 - \underline{\theta}) + p_2 - \underline{\theta} - \Delta, \\
\pi^{(1b)} &\triangleq \left(1 - \frac{(1-\nu)(p_2 - \underline{\theta})}{\Delta}\right)(p_1 - \underline{\theta}) + \nu(p_2 - \underline{\theta}), \\
\pi^{(2a)} &\triangleq p_1 - \underline{\theta} - \Delta + p_2 - \underline{\theta} - \Delta, \\
\pi^{(2b)} &\triangleq p_1 - \underline{\theta} - \Delta + \nu(p_2 - \underline{\theta}).
\end{aligned}$$

Pure strategy equilibria. We first analyze all possible equilibria in which all principals offer the same contract.

- Consider a candidate equilibrium in which all principals offer contract (2a), obtaining a profit $\pi^{(2a)}$, with $p_1 = p_2 = u'(1)$. From the analysis of the game with spot contracts, we already know that the semi-separating outcome cannot be achieved when $\Delta \leq \Delta_0$, which immediately implies that contract (1b) is not renegotiation-proof, as well as contract (2b), in the considered region of parameters. Thus, to establish the result, it only remains to be check whether a principal does not find it optimal to deviate offering contract (1a) — i.e., to shut down inefficient types only in the first period. The deviating principal would obtain a profit $\pi^{(1a)}$, and it can be immediately verified that, given the equilibrium prices, $\pi^{(1a)} < \pi^{(2a)}$ for all $\Delta \leq \Delta_0$.

- Consider a candidate equilibrium in which all principals offer contract (1a). In this case, each principal would obtain a profit $\pi^{(1a)}$, with $p_1 = u'(\nu)$ and $p_2 = u'(1)$. To see that this candidate equilibrium cannot exist, notice that, for all Δ, since $p_1 > p_2$, each principal has a profitable deviation consisting in offering contract (2b).

- By the same token, it can be easily checked that a candidate equilibrium in which all principals offer contract (2b) does not exists, since from $p_1 = u'(1) < p_2 = u'(\nu)$ it follows that any principal would find it optimal to deviate offering contract (1a).

- Consider a candidate equilibrium in which all principals offer contract (1b), and assume that this offer induces the semi-separating outcome — i.e., efficient agents accept the first period offer with probability γ^* defined by (3.2) and always produce in the second period. In this case, the equilibrium profit is $\pi^{(1b)}$, with $p_1 = u'(\nu\gamma^*)$ and $p_2 = u'(\nu)$. From the analysis of the game with spot contracts, we already know that if a principal deviates to contract (2b),[35] she would obtain a higher profit for $\Delta < \Delta_3$. Thus, to see whether, for every $\Delta \geq \Delta_3$, the considered candidate equilibrium exists, it remains to be checked whether, in this region of parameters, a principal does not find it optimal to deviate offering contract (1a).[36] However, $\pi^{(1a)} < \pi^{(2b)}$ whenever $p_1 > p_2$, which is the case in this equilibrium, implying that, for all $\Delta \geq \Delta_3$, there are no profitable deviations.

Mixed strategy equilibria. We now turn to consider all possible mixed strategy equilibria.

- Consider a candidate equilibrium in which each principal randomizes between offering contract (1a) and contract (2a). It can be immediately seen that, since $p_2 = u'(1)$, each principal would find it optimal to deviate offering contract (2b) for all $\Delta > \Delta_0$, whereas, for every $\Delta \leq \Delta_0$, principals are strictly better off offering contract (2a) than contract (1a) — i.e., there not exists a randomization which makes the principals indifferent between offering the two considered contract. We can thus conclude that the considered equilibrium does not exist.

- Consider a candidate equilibrium in which each principal randomizes between offering contract (1a) and contract (1b). In this candidate equilibrium, $p_1 > p_2$, which implies $\pi^{(2b)} > \pi^{(1a)}$ — i.e., each principal has a profitable deviation consisting in offering contract (2b). This proves that the considered equilibrium does not exist.

- Consider a candidate equilibrium in which each principal offers contract (1a) with probability β and contract (2b) otherwise. In this candidate equilibrium, market prices are $p_1 = u'(\beta\nu + 1 - \beta)$ and $p_2 = u'(\beta + (1-\beta)\nu)$. Imposing $\pi^{(1a)} = \pi^{(2b)}$ immediately yields $\beta^* = \frac{1}{2}$, implying that this equilibrium can exists for $\Delta \geq \Delta^*$, since contract (2b) is renegotiation-proof if $\Delta \geq (1-\nu)(p_2 - \underline{\theta})$. The only possibly profitable deviation for each principal consists in offering contract (1b) — i.e., to induce the semi-separating outcome.[37]. After some algebra, the deviation profit can be written as $\pi^{(1b)} \triangleq \frac{\Delta^*}{1-\nu}\left(1 + \nu - \frac{\Delta^*}{\Delta}\right)$, which turns out to be lower than the equilibrium profit for every $\Delta < \frac{\Delta^*}{\sqrt{1-\nu}}$.[38]

[35] Notice that, under the necessary condition (3.9) for the attainability of a semi-separating outcome (which, in this case, gives $\Delta \geq \Delta_1$), it follows that contract (2b) is renegotiation-proof.

[36] This is the case since, in the considered region of parameters, it can be immediately checked that the deviation consisting in offering contract (2a) is never profitable.

[37] In fact, it can be immediately verified that, in the considered region of parameters, offering contract (2a) never constitutes a profitable deviation.

[38] To see this, notice that, for $\beta^* = \frac{1}{2}$, the equilibrium profit can be rewritten as $\frac{1+\nu}{1-\nu}\Delta^* - \Delta$.

- Consider a candidate equilibrium in which each principal randomizes between offering contract (2b) and contract (1b), which are renegotiation-proof for $\Delta > \Delta_1$ (since $p_2 = u'(\nu)$). Denoting by ρ the probability with which principals offer contract (1b), and imposing that principals are indifferent between offering these two contracts yields
$$\Delta^2 = \Delta_1(u'(\rho\nu\gamma^* + 1 - \rho) - \underline{\theta}),$$
where γ^* is given by (3.2), from which we obtain the equilibrium value ρ^*. Moreover, $\rho^* < 1$ if $\Delta < \Delta_3$. However, in order for this candidate equilibrium to exist, any principal must not find it profitable to deviate offering contract (1a). Imposing this condition yields $\Delta > \hat{\Delta}$.

- Consider a candidate equilibrium in which each principal randomizes between offering contract (2a) and contract (2b). It can be immediately seen that this equilibrium does not exist. Indeed, from the analysis of the static contract, we know that principals, who all offered $w_1 = \overline{\theta}$, are indifferent between offering $w_2 = \underline{\theta}$ and $w_2 = \overline{\theta}$ if $\Delta \in (\Delta_0, \Delta_1)$. However, in this region of parameters, it can be immediately seen that any principal has incentive to deviate by offering contract (1a).

- Consider a candidate equilibrium in which each principal randomizes between offering contract (2a) and contract (1b). It can be immediately seen that this equilibrium does not exist. In fact, contract (1b) is renegotiation proof if and only if $\Delta > (1-\nu)(p_2 - \underline{\theta})$. However, under the same condition, from the analysis of the static game, it follows that principals are strictly better off offering contract (2b) than contract (2a). For the same reason, also candidate equilibria in which each principal randomizes among offering contracts (1a), (2a) and (1b), or offering the four considered long-term renegotiation-proof contracts cannot exist.

- By a similar reasoning, also an equilibrium in which each principal randomizes among contracts (1b), (2a) and (2b) cannot exist. In fact, in order for this equilibrium to exist, principals who offered $w_1 = \overline{\theta}$ (i.e., contracts (2a) and (2b)) must be indifferent between offering $w_2 = \underline{\theta}$ and $w_2 = \overline{\theta}$, which is the case if $\Delta = (1-\nu)(p_2 - \underline{\theta})$. However, if this equality is fulfilled, then contract (1b) is not renegotiation-proof.

- Consider a candidate equilibrium in which each principal offers contract (1a) with probability β, contract (2a) with probability ϵ and (2b) otherwise. In this candidate equilibrium, market prices are $p_1 = u'(1-(1-\nu)\beta)$ and $p_2 = u'(1-(1-\nu)(1-\beta-\epsilon))$. By imposing $\pi^{(1a)} = \pi^{(2a)}$ and $\pi^{(2a)} = \pi^{(2b)}$, it is easy to see that it must be $p_1 =$

$p_2 \triangleq p$ and $\Delta = (1-\nu)(p-\underline{\theta})$,[39] from which it is trivial to find that $\beta^* = \alpha^*$ defined by (3.1) and $\epsilon^* = 1 - 2\alpha^*$. This constitutes an equilibrium whenever $\alpha^* \in (0, \frac{1}{2})$, yielding $\Delta \in (\Delta_0, \Delta^*)$.[40]

- Finally, consider a candidate equilibrium in which each principal offers contract (1a) with probability ϕ, contract (1b) with probability φ and (2b) otherwise. Hence, $p_1 = u'(\phi\nu + \varphi\nu\gamma^* + 1 - \phi - \varphi)$ and $p_2 = u'(\phi + (1-\phi)\nu)$. By imposing $\pi^{(1a)} = \pi^{(2b)}$ and $\pi^{(1b)} = \pi^{(2b)}$ we obtain $p_1 = p_2 \triangleq p$ and $\Delta^2 = (1-\nu)(p-\underline{\theta})^2$, from which we pin down the equilibrium values ϕ^* and φ^*. In order for this equilibrium to exist, the following conditions must hold:[41]

$$\phi^* \in (0,1) \iff \Delta \in \left(\frac{\Delta_0}{\sqrt{1-\nu}}, \frac{\Delta_1}{\sqrt{1-\nu}}\right),$$

$$\varphi^* \in (0,1) \iff \Delta \in \left(\frac{\Delta^*}{\sqrt{1-\nu}}, \sqrt{1-\nu}\left(u'\left(\frac{1+\nu-\sqrt{1-\nu}}{2}\right) - \underline{\theta}\right)\right),$$

and

$$\varphi^* + \phi^* < 1 \iff \Delta < \sqrt{1-\nu}\left(u'\left(\frac{(1+\nu)\sqrt{1-\nu}-1}{2\sqrt{1-\nu}-1}\right) - \underline{\theta}\right).$$

After some algebra, these conditions yield $\Delta \in \left(\frac{\Delta^*}{\sqrt{1-\nu}}, \hat{\Delta}\right)$. ■

Thus, for every $\Delta \leq \Delta_0$ and $\Delta \geq \hat{\Delta}$, the equilibrium with long-term renegotiation-proof contracts coincides with the equilibrium with spot contracts. The reason is rather simple. When the adverse selection problem is weak ($\Delta \leq \Delta_0$), principals never it profitable to shut down inefficient agents in first period — i.e., to offer the long-term renegotiation proof contract (1a) — and, as in the game with spot contracts, in the equilibrium, all agents produce in both periods. When the adverse selection problem is too severe ($\Delta \geq \hat{\Delta}$), in equilibrium, quantities are increasing over time (see Proposition 3.5), hence $p_1^* > p_2^*$. This in turn implies that offering contract (2b) (i.e., shutting down the inefficient type in the second period) is more appealing than offering contract (1a), in which the inefficient type is shut down in the first period, when the equilibrium price is higher. As a consequence, for all $\Delta \geq \hat{\Delta}$, the feasibility of a fully separating outcome in the first period does not alter the equilibrium set of the game as compared with the case with spot contracts.
By contrast, when the severity of the adverse selection problem takes intermediate val-

[39] The latter equality implies that contract (2b) is renegotiation-proof.

[40] In fact, there are no possibly profitable deviations to examine, since contract (1b) can be easily seen to be not renegotiation-proof in the considered region of parameters.

[41] These conditions are also sufficient for the existence of the considered equilibrium, since there are no possibly profitable deviations to examine, given that offering contract (2a) is never optimal in the region of parameters in which contract (1b) is renegotiation proof.

ues (i.e., for all $\Delta \in (\Delta_0, \hat{\Delta})$), the larger number of available outcomes under long-term renegotiation-proof contracts completely change the equilibrium set of the game. Specifically, there are three mixed strategy equilibria and, in all them, contract (1a), which corresponds to the separating outcome of the game with spot contracts, is offered with positive probability. These equilibria have the following features.

Proposition 3.B.2. *For all $\Delta \leq \hat{\Delta}$, aggregate output is constant over time.*

Proof. Aggregate quantities are as follows:

- for $\Delta \leq \Delta_0 : y_1^* = y_2^* = 1$;
- for $\Delta \in (\Delta_0, \Delta^*] : y_1^* = y_2^* = u'^{-1}\left(\frac{\Delta}{1-\nu} + \underline{\theta}\right)$;
- for $\Delta \in \left(\Delta^*, \frac{\Delta^*}{\sqrt{1-\nu}}\right] : y_1^* = y_2^* = \frac{1+\nu}{2}$;
- for $\Delta \in \left(\frac{\Delta^*}{\sqrt{1-\nu}}, \hat{\Delta}\right] : y_1^* = y_2^* = u'^{-1}\left(\frac{\Delta}{\sqrt{1-\nu}} + \underline{\theta}\right)$;
- for $\Delta \in \left(\hat{\Delta}, \Delta_3\right] : y_1^* = u'^{-1}\left(\frac{\Delta^2}{\Delta_1} + \underline{\theta}\right) < y_2^* = \nu$;
- for $\Delta > \Delta_3 : y_1^* = \nu\gamma^* < y_2^* = \nu$. ■

Proposition 3.B.3. *When mixed strategy equilibria are played,*

- *principals' profit is non monotone with respect to Δ: it is decreasing in Δ for $\Delta \in \left(\Delta^*, \frac{\Delta^*}{\sqrt{1-\nu}}\right)$ and decreasing otherwise;*
- *consumer welfare is a non-increasing function of Δ: it is constant for $\Delta \in \left(\Delta^*, \frac{\Delta^*}{\sqrt{1-\nu}}\right)$ and strictly decreasing otherwise.*

Proof. Principals' profits in the mixed strategy equilibria are as follows:

- for $\Delta \in (\Delta_0, \Delta^*] : \pi^* = \frac{2\nu}{1-\nu}\Delta$;
- for $\Delta \in \left(\Delta^*, \frac{\Delta^*}{\sqrt{1-\nu}}\right] : \pi^* = \frac{1+\nu}{1-\nu}\Delta^* - \Delta$;
- for $\Delta \in \left(\frac{\Delta^*}{\sqrt{1-\nu}}, \hat{\Delta}\right] : \pi^* = \left(\frac{1+\nu}{\sqrt{1-\nu}} - 1\right)\Delta$;
- for $\Delta \in \left(\hat{\Delta}, \Delta_3\right] : \pi^* = \frac{\Delta^2}{\Delta_1} - \Delta + \frac{\nu}{1-\nu}\Delta_1$.

As for consumer welfare, for any $\Delta \in \left(\Delta_0, \hat{\Delta}\right)$ quantity is constant over time and non-increasing in Δ, implying that also consumer welfare is non-increasing in Δ. Finally, for $\Delta \in \left(\hat{\Delta}, \Delta_3\right]$ the equilibrium outcome is as in the model with spot contracts, and we already proved that consumer welfare is decreasing in Δ. ■

Finally, in the game with discounted payoffs, we can prove what follows.

Proposition 3.B.4. *For all $\delta \in (0,1)$, the equilibrium outcome with spot contracts coincides with the equilibrium outcome with long-term contracts and Pareto-improving renegotiations.*

Proof. Long-term renegotiation-proof contracts are as in the game with no discounting. However, when principals discount future profits, a fully separating outcome can be achieved also with spot contracts. Therefore, there is a one-to-one correspondence among the outcomes (i.e., discounted wages and market prices) that can be achieved with long-term renegotiation-proof contracts and with short-term contracts, which establishes the result. Accordingly, the mixed strategy equilibria of the game with spot contracts can be equivalently regarded as randomizations among the corresponding long-term renegotiation-proof contracts. Specifically, the mixed strategy equilibrium for $\Delta \in (\Delta_0, \Delta^*]$ can be equivalently obtained through randomization among three long-term renegotiation-proof contracts:

(i) with probability α^* defined in (3.1), each principal offers contract (1a).

(ii) with the same probability α^*, each principal offers contract (2b). Such contract is renegotiation-proof if and only if the principals after the first period would not be better off by offering $w_2 = \overline{\theta}$, that is for $\Delta \geq (1-\nu)(p_2^* - \underline{\theta})$, which is satisfied (as an equality).

(iii) with complementary probability $1 - 2\alpha^*$, each principal offers contract (2a).

Analogously, the mixed strategy equilibrium for $\Delta \in (\Delta^*, \Delta_1)$ can be obtained if principal randomize between the following long-term renegotiation-proof contracts:

(i) with probability β^*, defined as the unique solution of equation (3.6), each principal offers contract (1a).

(ii) with complementary probability $1 - \beta^*$, each principal offers contract (2b).[42] ■

[42] Contract (ii) is renegotiation-proof for $\Delta \geq (1-\nu)(p_2^* - \underline{\theta}) = (1-\nu)(u'(\beta^* + \nu(1-\beta^*)) - \underline{\theta})$. Since the right-hand side is increasing in $\beta^* > \frac{1}{2}$ and the inequality is satisfied at $\beta^* = \frac{1}{2}$, we can conclude that the considered contract is renegotiation-proof.

Conclusion

The three essays composing this thesis were aimed at shading some light on the key economic forces behind the behavior of firms in different real-world contexts, and derive some policy implications. However, several other relevant aspects related to the topics covered in this thesis remains to be addressed. In this concluding section, it is worth outlining some of the directions for future research.

As for the role played by professional advisors in M&A transactions, besides its relevance in the merger review process, which has been investigated in Chapter 1, it may be interesting to analyze other potential conflicts of interest between the advisors and the hiring firm, for instance those arising when the advisor is a commercial bank involved also in financing the acquisition (see, e.g., Allen et al., 2004). In this case, the bank can advise a firm to undertake a merger which is not in the shareholders' interest but maximizes its own business volumes. Thus, an acquiring firm faces non trivial problems concerning how to structure an optimal advisory contract and, relatedly, how to finance the acquisition.

The evaluation of the pro- and anti-competitive effects of platform parity agreements has been carried out in Chapter 2 taking as given the industry structure. However, the dynamic effects of these clauses remains to be established. To be more specific, it would be interesting to understand how the presence of platform parity impacts on the incentives to pursue horizontal integrations among platforms or vertical integrations along the supply chain, and the role of parity provisions in facilitating or preventing firms' entry at different levels of the supply chain. Clearly, the consideration of these aspects is of crucial importance in the decisions of antitrust authorities concerning the prohibition of platform parity clauses.

In line with the previous literature, in the work examining the interplay between market competition and the ratchet effect, presented in Chapter 3, it is assumed that a principal can learn the agent's private information concerning the production cost only through incentive-compatible contracts. However, nowadays, data analytics can be employed in performance measurement systems, which implies that informational asymmetries may be overcome by acquiring information from a third-party, for instance a data analytics supplier (see, e.g., Kastl et al., 2018). However, whether and to what extent, in equilibrium, the availability of data analytics is effective in mitigating the ratchet effect remains to be established, since the answer crucially depends on the incentives of the data analytics firms to supply information of good quality at an affordable price. Moreover, clearly, both upstream competition (in the market of data analytics suppliers) and downstream competition (in the firms' output market) is likely to play a major role in shaping the equilibrium outcome.

Bibliography

Acharya, A., & Ortner, J. (2017). Progressive Learning. *Econometrica*, 85(6), 1965-1990.

Akman, P. (2016). A competition law assessment of platform most-favored-customer clauses. *Journal of Competition Law & Economics*, 12(4), 781-833.

Allen, L., Jagtiani, J., Peristiani, S., & Saunders, A. (2004). The Role of Bank Advisors in Mergers and Acquisitions. *Journal of Money, Credit, and Banking*, 36(2), 197-224.

Amir, R., Diamantoudi, E., & Xue, L. (2009). Merger performance under uncertain efficiency gains. *International Journal of Industrial Organization*, 27(2), 264-273.

Armstrong, M., & Vickers, J. (2010). A model of delegated project choice. *Econometrica*, 78(1), 213-244.

Avilés-Lucero, F., & Boik, A. (2018). Wholesale most-favored-nation clauses and price discrimination with negative consumption externalities: equivalence results. *Journal of Regulatory Economics*, 54(3), 266-291.

Baker, J. B., & Chevalier, J. A. (2012). The competitive consequences of most-favored-nation provisions. *Antitrust*, 27, 20-26.

Bao, J., & Edmans, A. (2011). Do investment banks matter for M&A returns?. *The Review of Financial Studies*, 24(7), 2286-2315.

Beccuti, J., & Möller, M. (2018). Dynamic adverse selection with a patient seller. *Journal of Economic Theory*, 173, 95-117.

Besanko, D., & Spulber, D. F. (1993). Contested mergers and equilibrium antitrust policy. *The Journal of Law, Economics, and Organization*, 9(1), 1-29.

Boik, A., & Corts, K. S. (2016). The effects of platform most-favored-nation clauses on competition and entry. *The Journal of Law and Economics*, 59(1), 105-134.

Bolton, P., & Dewatripont, M. (2005). *Contract theory.* MIT press.

Butz, D. A. (1990). Durable-Good Monopoly and Best-Price Provisions. *The American Economic Review*, 80(5), 1062-1076.

Caillaud, B., Jullien, B., & Picard, P. (1995). Competing Vertical Structures: Precommitment and Renegotiation. *Econometrica*, 63, 621-646.

Calzada, J., Manna, E., & Mantovani, A. (2019). Platform Price Parity Clauses and Segmentation. UB Economics-Working Papers, 2018, E19/387.

Calzolari, G., Denicolò, V., & Zanchettin, P. (2017). Exclusive dealing with distortionary pricing. Working paper.

Carmichael, H. L., & MacLeod, W. B. (2000). Worker cooperation and the ratchet effect. *Journal of Labor Economics*, 18(1), 1-19.

Cazaubiel, A., Cure, M., Johansen, B. O., & Vergé, T. (2018). Substitution Between Online Distribution Channels: Evidence from the Oslo Hotel Market. University of Bergen, Department of Economics (No. 8/18).

Choné, P., & Linnemer, L. (2008). Assessing horizontal mergers under uncertain efficiency gains. *International Journal of Industrial Organization*, 26(4), 913-929.

Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386-405.

Coase, R. H. (1972). Durability and monopoly. *The Journal of Law and Economics*, 15, 143-149.

Cooper, T. E. (1986). Most-favored-customer pricing and tacit collusion. *The RAND Journal of Economics*, 17(3), 377-388.

Cremer, J., & Riordan, M. H. (1987). On governing multilateral transactions with bilateral contracts. *The RAND Journal of Economics*, 436-451.

Cunha, M., Sarmento, P., & Vasconcelos, H. (2014). Uncertain efficiency gains and merger policy. Mimeo.

DeGraba, P. (1996). Most-Favored-Customer Clauses and Multilateral Contracting: When Nondiscrimination Implies Uniformity. *Journal of Economics & Management Strategy*, 5(4), 565-579.

DeGraba, P., & Postlewaite, A. (1992). Exclusivity clauses and best price policies in input markets. *Journal of Economics & Management Strategy*, 1(3), 423-454.

Deneckere, R., & Davidson, C. (1985). Incentives to form coalitions with Bertrand competition. *The RAND Journal of economics*, 16(4), 473-486.

Dewatripont, M. (1989). Renegotiation and information revelation over time: The case of optimal labor contracts. *The Quarterly Journal of Economics*, 104(3), 589-619.

Edelman, B., & Wright, J. (2015). Price coherence and excessive intermediation. *The Quarterly Journal of Economics*, 130(3), 1283-1328.

Ezrachi, A. (2015). The competitive effects of parity clauses on online commerce. *European Competition Journal*, 11(2-3), 488-519.

Farrell, J., & Shapiro, C. (1990). Horizontal mergers: an equilibrium analysis. *The American Economic Review*, 80(1), 107-126.

Foros, O., Kind, H. J., & Shaffer, G. (2017). Apple' s agency model and the role of most-favored-nation clauses. *The RAND Journal of Economics*, 48(3), 673-703.

Freixas, X., Guesnerie, R., & Tirole, J. (1985). Planning under incomplete information and the ratchet effect. *The Review of Economic Studies*, 52(2), 173-191.

Fudenberg, D., & Tirole, J. (1983). Sequential bargaining with incomplete information. *The Review of Economic Studies*,50(2), 221-247.

Fudenberg, D., & Tirole, J. (1991). *Game theory.* MIT Press.

Gaudin, G. (2019). Vertical relations, opportunism, and welfare. *The RAND Journal of Economics*, 50(2), 342-358.

Gerardi, D., & Maestri, L. (2018). Dynamic Contracting with Limited Commitment and the Ratchet Effect. Working Paper.

Gibbons, R. (1987). Piece-rate incentive schemes. *Journal of Labor Economics*, 5(4), 413-429.

Giebe, T., & Lee, M. (2019). Competitors In Merger Control: Shall They Be Merely Heard Or Also Listened To?. Available at SSRN 2569832.

Golubov, A., Petmezas, D., & Travlos, N. G. (2012). When it pays to pay your investment banker: New evidence on the role of financial advisors in M&As. *The Journal of Finance*, 67(1), 271-311.

Gromb, D., & Martimort, D. (2007). Collusion and the organization of delegated expertise. *Journal of Economic Theory*, 137(1), 271-299.

Grossman, S. J., & Hart, O. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4), 691-719.

Gul, F., Sonnenschein, H., & Wilson, R. (1986). Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39(1), 155-190.

Gupta, D., & Gerchak, Y. (2002). Quantifying operational synergies in a merger/acquisition. *Management Science*, 48(4), 517-533.

Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1), 1-23.

Hart, O. (1983). The Market Mechanism as an Incentive Scheme. *The Bell Journal of Economics,* 14, 366-382.

Hart, O., & Tirole, J. (1988). Contract renegotiation and Coasian dynamics. *The Review of Economic Studies,* 55(4), 509-540.

Hart, O., & Tirole, J. (1990). Vertical integration and market foreclosure. *Brookings papers on economic activity. Microeconomics*, 205-286.

Horn, H., & Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, 408-419.

Hosios, A. J., & Peters, M. (1993). Self-Selection and Monitoring in Dynamic Incentive Problems with Incomplete Contracts. *The Review of Economic Studies*, 60(1), 149-174.

Johansen, B. O., & Vergé, T. (2016). Platform price parity clauses with direct sales. IDEI working.

Johnson, J. P. (2017). The agency model and MFN clauses. *The Review of Economic Studies*, 84(3), 1151-1185.

Kanemoto, Y., & MacLeod, W.B. (1992). The ratchet effect and the market for second-hand workers. *Journal of Labor Economics*, 10(1), 85-98.

Kastl, J., Pagnozzi, M., & Piccolo, S. (2018). Selling information to competitive firms. *The RAND Journal of Economics*, 49(1), 254-282.

Katsoulacos, Y., & Ulph, D. (2009). On optimal legal standards for competition policy: a general welfare-based analysis. *The Journal of Industrial Economics*, 57(3), 410-437.

Laffont, J. J., & Tirole, J. (1987). Comparative statics of the optimal dynamic incentive contract. *European Economic Review*, 31(4), 901-926.

Laffont, J. J., & Tirole, J. (1988). The dynamics of incentive contracts. *Econometrica*, 1153-1175.

Laffont, J. J., & Tirole, J. (1990). Adverse selection and renegotiation in procurement. *The Review of Economic Studies*, 57(4), 597-625.

Legros, P., & Newman, A. (2013). A Price Theory of Vertical and Lateral Integration. *The Quarterly Journal of Economics*, 128(2), 725-770.

Lyons, B. R. (2003). Could politicians be more right than economists? A theory of merger standards. Mimeo.

Maestri, L. (2017). Dynamic Contracting under Adverse Selection and Renegotiation. *Journal of Economic Theory*, 171, 136-173.

Martin, S. (1993). Endogenous Firm Efficiency in a Cournot Principal-Agent Model. *Journal of Economic Theory,* 59, 445-450.

Marx, L.M., & Shaffer, G. (2004). Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity: Comment. *The American Economic Review*, 94(3), 796-801.

McAfee, R. P., & Schwartz, M. (1994). Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *The American Economic Review*, 210-230.

McAfee, R. P., & Williams, M. A. (1992). Horizontal mergers and antitrust policy. *The Journal of Industrial Economics*, 40(2), 181-187.

Motta, M. (2004). *Competition policy: theory and practice*. Cambridge University Press.

Motta, M., & Tarantino, E., (2018). The Effect of horizontal mergers, when firms compete in prices and investments. Mimeo.

Nocke, V., & Whinston, M. D. (2010). Dynamic merger review. *Journal of Political Economy*, 118(6), 1200-1251.

Nocke, V., & Whinston, M. D. (2013). Merger policy with merger choice. *The American Economic Review*, 103(2), 1006-1033.

Ortner, J. (2017). Durable goods monopoly with stochastic costs. *Theoretical Economics*, 12(2), 817-861.

Pagnozzi, M., & Piccolo, S. (2011). Vertical separation with private contracts. *The Economic Journal*, 122(559), 173-207.

Raith, M. (2003). Competition, Risk and Managerial Incentives. *American Economic Review,* 93, 1425-1436.

Rau, P. R. (2000). Investment bank market share, contingent fee payments, and the performance of acquiring firms. *Journal of Financial Economics*, 56(2), 293-324.

Rey, P., & Tirole, J. (2007). A primer on foreclosure. Handbook of industrial organization, 3, 2145-2220.

Rey, P., & Vergé, T. (2004). Bilateral control with vertical contracts. *The RAND Journal of Economics*, 35(4), 728-746.

Rey, P., & Vergé, T. (2017). Secret contracting in multilateral relations. Working paper.

Ronayne, D., & Taylor, G. (2018). Competing Sales Channels. University of Oxford, Department of Economics (No. 843).

Sørgard, L. (2009). Optimal merger policy: enforcementvs. deterrence. *The Journal of Industrial Economics*, 57(3), 438-456.

Scharfstein, D. (1988). Product-market competition and managerial slack. *The RAND Journal of Economics*, 147-155.

Schnitzer, M. (1994). Dynamic duopoly with best-price clauses. *The RAND Journal of Economics*, 25(1), 186-196.

Servaes, H., & Zenner, M. (1996). The role of investment banks in acquisitions. *The Review of Financial Studies*, 9(3), 787-815.

Shen, B., & Wright, J. (2019). Why (don' t) firms free ride on an intermediary' s advice?. *International Journal of Industrial Organization*, 64, 27-54.

Simon, H. A. (1951). A formal theory of the employment relationship. *Econometrica*, 293-305.

Skreta, V. (2006). Sequentially optimal mechanisms. *The Review of Economic Studies*, 73(4), 1085-1111.

Skreta, V. (2015). Optimal auction design under non-commitment. *Journal of Economic Theory*, 159, 854-890.

Stokey, N. L. (1981). Rational expectations and durable goods pricing. *The Bell Journal of Economics*, 112-128.

Strulovici, B. (2017). Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts. *Econometrica*, 85, 585-616.

Tirole, J. (1988). *The theory of industrial organization.* MIT press.

Vasu, R. (2019). Optimal M&A advisory contracts. Mimeo.

Wang, C., & Wright, J. (2016). Search platforms: Showrooming and price parity clauses. Working Paper, National University of Singapore.

Whinston, M. D. (2007). Antitrust policy toward horizontal mergers. *Handbook of industrial organization*, 3, 2369-2440.

Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *The American Economic Review*, 58(1), 18-36.

Williamson, O. E. (1973). Markets and hierarchies: some elementary considerations. *The American Economic Review*, 63(2), 316-325.