Controlling versus enabling*

Andrei Hagiu[†] and Julian Wright[‡]

September, 2017

Abstract

How does a firm decide whether to employ professionals and control how they deliver services to clients, or to operate as a platform enabling independent professionals to provide services directly to clients? Similarly, how does a manufacturer decide whether to allow sales agents to choose certain costly actions (e.g. kickbacks to clients) or to take control of these actions itself? We answer this question using a principal-agent framework in which both the principal and the agent must be incentivized to carry out investments (or effort) that increase the revenue they jointly create. Our theory explains when the principal should take control over a particular decision ("control") or should instead allow the agent to make the decision ("enable"). It does so both for the case when there are multiple such transferable decisions for a single agent, and for the case when there are many agents and one transferable decision for each. We also consider the possibility of cost asymmetries between the principal and the agent, spillovers across agents, and the misclassification of the principal as an employer even though agents are allocated the relevant control rights. Finally, we explain how the "control vs. enable" choice and its associated tradeoffs differ from the classic "make vs. buy" choice.

JEL classification: D23, D86, L14

Keywords: control rights, platforms, employment, independent contractors.

*We thank Oliver Hart, Justin Johnson, Bruno Jullien, Marc Rysman, Steve Tadelis, Jean Tirole, Eric Van den Steen, Birger Wernerfelt, participants in the Platform Strategy Research Symposium at Boston University, the Platforms Workshop at the National University of Singapore, the Conferences on the Economics of Intellectual Property, Software and the Internet and Advances in the Economics of Organization and the New Economy in Honor of Jacques Crémer at the Toulouse School of Economics, and participants in seminars at Boston University, the Federal Trade Commission, Georgetown University, Harvard University, MIT, the National University of Singapore, Northwestern University, University of California San Diego, University of Southern California, and University of Utah, as well as the Department Editor and two anonymous referees for their helpful comments. Andy Lee provided excellent research assistance. Julian Wright gratefully acknowledges research funding from the Singapore Ministry of Education Academic Research Fund Tier 1 Grant No. R122000215-112.

†MIT Sloan School of Management, Cambridge, MA 02142, E-mail: ahagiu@mit.edu

‡Department of Economics, National University of Singapore, Singapore 117570, E-mail: jwright@nus.edu.sg

1 Introduction

A key decision for many firms is whether to control the provision of services to customers by employing workers or whether to enable independent contractors to take control of service provision. This decision has been relevant in some industries for a long time—such as manufacturers and sales agents. However, it has become more prominent in recent times, reflecting that in a rapidly increasing number of service industries (e.g. consulting, education, home services, legal, outsourcing, staffing, and taxi services), online platforms have emerged to take advantage of information and remote collaboration technologies to enable independent professionals to directly connect with customers (e.g. Catalant, Coursera, Gerson Lehrman Group, Task Rabbit, Uber, and Upwork). These firms typically differ from their more traditional counterparts by letting professionals control some or all of the relevant decision rights, such as prices, expenditure on the quality and maintenance of equipment, and marketing of the professionals. This contrast motivates our theoretical study of a principal's choice of whether to keep decision authority or to grant it to an agent.

In these settings, the revenues generated by the principal typically depend on both its ongoing investments (e.g. in the quality of its product or technological infrastructure) and those made by the agent (e.g. the effort the agent puts into providing a high quality service). When neither the principal's nor the agent's investments are contractible, joint production calls for some sharing of revenues between the principal and the agent to ensure each has an incentive to invest. At the same time, as noted above, there are other non-contractible decisions which also affect revenues but can be controlled by either the principal or the agent. In this paper, we study the optimal allocation of control rights over these *transferable* decision variables, taking into account that revenue sharing affects all decisions.

To do so, we develop a model that contains three types of non-contractible decisions: two costly and non-transferable investment decisions—one for the principal and one for the agent—and a set of transferable decisions, each of which can be controlled by either the principal or the agent. Our analysis yields three main sets of results.

First, we show that when there is a single agent, multiple transferable decisions and sufficiently small (or no) cost differences between the principal and the agent in undertaking transferable actions, control rights over these decisions should all be given to the same party (principal or agent), namely the party that obtains the higher revenue share in equilibrium. In other words, low-powered incentives (i.e. control over the transferable decisions) should be aligned with high-powered incentives (i.e. higher share of revenues) in order to minimize revenue-sharing distortions. An implication of this result is that in dealing with a single agent, the principal only has to choose between two modes of organization: what we call the $\mathcal{P}$-mode, in which the principal keeps control over all transferable decisions (this is the "control" mode, which can be interpreted as employment), and what we call the $\mathcal{A}$-mode, in which the principal gives control over all transferable decisions to the agent (this is the "enable" mode, which can be interpreted as independent contracting[1]). This result does not rely on any interaction

[1] The reason for using the term "enable" is that, although the principal gives the agent control over all transferable decisions, the principal still controls its own non-transferable decision. This decision can be interpreted as investment in

effects among the various decisions in the revenue function, or on any cost economies of scope across transferable decisions.

Second, we show that when the principal deals with multiple agents, a mixed mode of organization can be optimal, in which some agents are given control over transferable decisions while others are not. This result holds even though all agents are identical, there are no spillovers across the transferable decisions of different agents, and there are no diseconomies of scale in dealing with agents in the same mode. Given the platform nature of the principal's non-transferable investment (i.e. that investment enhances revenues generated by all agents), a mixed mode across agents becomes a strategic way for the principal to get some of the advantages of each of the two pure modes, and do better than both.

Third, we study the choice between $\mathcal{P}$-mode and $\mathcal{A}$-mode when there are multiple agents and spillovers across the transferable decisions of different agents. In this case, the spillover-induced distortion shifts the baseline tradeoff between $\mathcal{P}$-mode and $\mathcal{A}$-mode by either exacerbating the revenue-sharing distortion (which favors the $\mathcal{P}$-mode) or offsetting it (which favors the $\mathcal{A}$-mode). The latter scenario leads to some counter-intuitive results, particularly when compared to the corresponding classic "make vs. buy" predictions.

For example, consider the case in which the transferable decision is a revenue-increasing, costly investment (e.g. giving kickbacks to clients) and the spillovers are negative (e.g. a sales agent of a given manufacturer steals business from the manufacturer's other sales agents by giving clients greater kickbacks). In $\mathcal{A}$-mode, individual agents invest too much in such kickbacks by not fully internalizing the spillovers. However, this can help offset the revenue-sharing distortion, namely that the party with control rights invests too little in kickbacks because it keeps less than 100% of the revenue generated. Thus, the $\mathcal{A}$-mode can be a useful way for the principal to get agents to choose higher levels of the transferable decision variable without giving them an excessively high share of revenues. This mechanism has two counterintuitive consequences. First, when negative spillovers are not too large in magnitude, an increase in their magnitude shifts the tradeoff in favor of the $\mathcal{A}$-mode—the opposite of the standard "internalize externalities" logic. Second, if the magnitude of negative spillovers is sufficiently large, the $\mathcal{A}$-mode (respectively, the $\mathcal{P}$-mode) is more likely to be chosen when the principal's (respectively, the agents') moral hazard becomes more important. This is the opposite of the standard "give control to the party whose investments are more important" prediction, which prevails when spillovers are positive and which is found in the classic make vs. buy literature. We extend these results with spillovers to the case when the transferable decision variable is price and show that similar counter-intuitive results can prevail in this case too.

Our theory provides a natural way to conceptualize the fundamental difference between traditional firms that hire employees and platforms that enable independent contractors to interact with customers, based on the allocation of control rights between the firm and workers over decisions that affect the revenues generated from customers. Simply put, firms that allocate more control rights to workers are closer to the platform/marketplace model. In light of this, we also explore what happens when the firm's choice of mode is misclassified by regulators, in particular when the firm is required

an infrastructure (or platform) that enables the agent to interact with customers.

to provide employment benefits to an agent even if the agent holds full control rights over transferable decisions. While such a misclassification has no effect on the outcome in our benchmark model, we show that if the agent faces a liquidity constraint, it can cause the principal to recover profits inefficiently through a higher share of variable revenues, resulting in lower profits and lower welfare.

2 Related literature

At a high level, a key contribution of our paper is to provide a way of extending the classic "make vs. buy" literature in strategy and economics to incorporate the study of platforms (the "enable" mode). We defer the discussion of how the "control vs. enable" choice and its associated tradeoffs differ from the classic "make vs. buy" choice to Section 7, because it also provides a way to more clearly articulate the strategy implications of our results.

Our model also relates to the literature on decision authority within organizations (e.g. Aghion and Tirole, 1997, Alonso et al., 2008, Bester, 2009). Specifically, one can view the principal in our model as an owner and manager of a firm deciding on the allocation of decision rights between herself and an employee (and on the corresponding compensation structure), in a context in which both parties must make on-going investments that affect total revenues generated. A key difference relative to this literature is that in our model the misalignment of objectives between principal and agent is endogenously determined by the principal's choice of revenue sharing (which in turn is driven by the underlying double-sided moral hazard problem). By contrast, most of this literature typically assumes exogenously given misalignments of objectives between the various parties involved. Furthermore, the driving forces in our model are the distortions due to double-sided moral hazard and spillovers, whereas the existing literature on decision authority focuses on uncertainty, information asymmetries and cheap talk.

Since in our model revenues must be shared between the principal and the agent to incentivize both sides to make non-contractible investments, we directly build upon principal-agent models with double-sided moral hazard (Romano, 1994, Bhattacharyya and Lafontaine, 1995). The key difference relative to these papers is that we introduce a third type of non-contractible decision, control over which can be allocated to either the principal or the agents. We also generalize these settings by allowing each type of decision variable to be multi-dimensional.

At a high level, our explanation for why a mixed mode across agents may be optimal is related to that provided by Bai and Tao (2000) for why some business format franchisors choose to operate with a mix of company-owned units and franchised units. There is, however, an important difference: Bai and Tao's explanation relies on the public good problem created by the fact that each franchise manager (agent) makes an independent investment in goodwill, which increases revenues of other franchisees. Since managers must also invest in sales efforts for their own unit (unlike goodwill, these efforts generate no spillovers), using a mixed mode allows the franchisor (principal) to ensure sufficient goodwill by having a positive number of managers under low-powered contracts (our $\mathcal{P}$-mode), as well as sufficient sales effort by having a positive number of managers under high-powered contracts (our

$\mathcal{A}$-mode). By contrast, in our set-up in Section 5.1 (where we focus on the possible optimality of mixed modes across agents), there is no public good problem: the principal's non-transferable investment is always chosen by the principal, and the transferable actions do not generate any spillovers across agents. Consequently, the result that the principal may choose a mixed mode is more surprising in our context. The underlying mechanism is also different.

Finally, our paper relates to a recently emerging strand of literature that studies conditions under which retailers/platforms take control over transferable decisions pertaining to the sale of products to end-consumers or allow their suppliers/complementors to keep control over these decisions. Bhardwaj (2001), Foros et al. (2013), Gans (2012), Abhishek et al. (2015), and Johnson (2017) focus on price as the main decision that can be controlled by the retailers/platforms (wholesale model) or by the suppliers/complementors (agency model). Desiraju and Moorthy (1997), Jerath and Zhang (2010), and Hagiu and Wright (2015a), study delegation of both price and costly investment (e.g. service) decisions. The key ingredient that makes our model very different from the models in these articles is the double-sided moral hazard underlying the delegation decision. We allow that both the principal (e.g. retailer or a platform in their settings) and the agent (e.g. a supplier or a platform complementor in their settings) make costly investments that affect realized demand and revenues. This is a key driver of our results: if, for example, only the agent faced a costly investment decision, in our setting control would always be given to the agent.

3 Examples

There are many industries in which the choice that we study is relevant. An important set of industries involves firms that can either employ professionals and control how they deliver services to clients, or operate as marketplaces enabling independent professionals to provide services directly to clients. This choice is relevant to both Internet-based service platforms (e.g. Catalant, Coursera, Handy, Lyft and Uber, Rubicon Global, Task Rabbit, and Upwork) and to firms operating in a number of "offline" industries.

The hair salon industry is a good example, as it has long featured two modes of organization, that can be viewed as corresponding to our $\mathcal{P}$-mode and $\mathcal{A}$-mode respectively. Some salons employ their hairstylists and pay them fixed hourly wages plus commissions that are a percentage of sales. Such salons control how individual hair dressers are promoted, provide most of the supplies and equipment that stylists use for hair cutting and styling, and determine prices ($\mathcal{P}$-mode). In contrast, other salons rent out chairs (booths) to independent hairstylists. The stylists keep all earnings minus fixed monthly booth rental fees that are paid to the salon. In such salons, individual hair stylists promote themselves, are responsible for providing and maintaining the majority of the supplies and equipment they need, and choose their prices individually ($\mathcal{A}$-mode). In both modes, the salon owners still make all necessary investments to maintain the facilities and advertise the salon to customers, while the stylists must exert effort to provide quality service to customers.

Another large set of relevant industries involves firms that need salespeople, brokers, or distributors to sell their products or services. Examples include the use of salespeople by manufacturers and the use of brokers by insurance companies. Firms in these markets often use a mix of independent agents, who among other things determine the extent of kickbacks they offer to purchase managers, and employees, for whom the firm determines and provides the kickbacks that are given to purchase managers. The commission rates paid out by the firms vary substantially across the two modes (Anderson, 1985).

Similarly, firms providing a wide range of products or services can do so through company-owned outlets or through independent franchisees. Most business format franchisors (e.g. hotels, fast-food outlets, and car rentals) use a combination of upfront fixed franchise fees and sales-based royalties (Blair and Lafontaine, 2005). While franchise contracts are notoriously restrictive, franchisees nevertheless control some key decisions that impact the revenues they generate (e.g. their expenditure on staff). In contrast, these decisions are made by the firm in company-owned outlets.

Table 1: Examples

	Transferable decisions	*Non-transferable investment decisions made by agents*	*Non-transferable investment decisions made by the principal*
Hair salons	marketing of individual hair dressers; quality and maintenance of equipment; spending on supplies	quality of service	maintenance and advertising of the salon
Transportation (e.g. Uber vs. traditional taxi companies)	quality of the car (make and model); maintenance of the car; location of work	knowledge of routes in the relevant area; quality of service	quality & maintenance of the technological infrastructure (payment, dispatch system); advertising of the firm
Consulting (e.g. Hourly Nerd vs. McKinsey) and outsourcing (e.g. Upwork vs. Infosys)	marketing of individual professionals and their skills; price	effort to understand customer requests; quality of service	quality & maintenance of the (online) system for communication, monitoring and payment; advertising of the firm
Online education (e.g. Coursera vs. University of Phoenix)	quality of the course design; advertising of individual instructors and courses	course preparation; quality of course delivery	quality & maintenance of the online infrastructure; advertising of the site
Waste and recycling (e.g. Rubicon Global vs. Waste Management)	condition and maintenance of equipment for waste collection and hauling	quality of service	quality & maintenance of the technological infrastructure (payment, scheduling); advertising of the firm
Producers and sales agents	kickbacks to clients	knowledge of product; sales effort	advertising of the product; product support
Franchising	expenditure on staff and their benefits	outlet manager effort	quality of the product; advertising of the brand
Sharecropping	quality of inputs (seeds, fertilizer and pesticides); tools and equipment ; bribes	adoption of high-yield farming practices; effort in working the land	large investments (e.g. maintenance of irrigation system)

An example that is more relevant for developing countries is sharecropping, in which landowners can decide how much to share their crops and relevant decision rights with agricultural workers. At one extreme, the landowner rents the land to a lessee at a fixed rate and the lessee has full control over inputs. At the other extreme, the landowner employs agricultural laborers at fixed wages and fully controls inputs. In between these two extremes, the landowner and the sharecropper share crops[2] and decision rights over inputs. Double-sided moral hazard is key in explaining the structure of sharecropping contracts, as noted by Bhattacharyya and Lafontaine (1995).

Table 1 shows how these and other examples fit our theory. In particular, it illustrates the three different types of non-contractible decision variables featured in our model that affect the revenue generated by each agent: (i) costly transferable decisions that are chosen by the principal in $\mathcal{P}$-mode and by the agent in $\mathcal{A}$-mode; (ii) costly ongoing investments always chosen by the agent; and (iii) costly ongoing investments always chosen by the principal.

Another possible non-contractible transferable decision variable is the price charged to customers. We consider this case in Section 6.3. However, in other cases, the price may be contractible, and indeed set by the principal in its contract with the agent. That possibility is easily handled, and we note below how our results extend to allow for this possibility. The price may also be pinned down by market constraints, in which case it can be treated as a fixed constant in our analysis (e.g. this may arise for some of the low-skill services offered through platforms such as Task Rabbit and Upwork, as well as for sharecropping).

4 One agent and multiple transferable actions

In this section, we analyze a setting in which there is a principal (e.g. a firm) and a single agent, but potentially many transferable actions. The revenue generated jointly by the principal and the agent if the latter accepts the principal's contract is a function of three types of actions, $a^1, .., a^M$, q and Q, all of which are non-contractible and are explained below. We assume the revenue function is linear in these actions:

$$R\left(a^1, .., a^M, q, Q\right) = \sum_{i=1}^{M} \beta^i a^i + \phi q + \Phi Q.$$

The actions q and Q are assumed to be non-transferable. Specifically, the agent always chooses $q \in \mathbb{R}_+$ at cost $c\left(q\right) = \frac{1}{2}q^2$ and the principal always chooses $Q \in \mathbb{R}_+$ at cost $C\left(Q\right) = \frac{1}{2}Q^2$. This means there is double-sided moral hazard: q encompasses ongoing effort and investment decisions that are always made by the agent and that raise the customers' willingness to pay for the service provided (see column 3 in Table 1), while Q captures the ongoing investments that are always made by the principal (see column 4 in Table 1). In contrast, the actions $a^1, .., a^M$ are all transferable, i.e. each of them can be chosen *either* by the principal or by the agent, depending on how the principal chooses to allocate control rights (see column 2 of Table 1). We assume $M \geq 1$ and define the vector $a \equiv \left(a^1, ..., a^M\right)$.[3]

[2] While 50/50 crop sharing is the most common practice, other splits are also used, as documented by Terpstra (1998).

[3] If $M = 0$, then we have a double-sided moral hazard problem without any transferable action to allocate control to. This has already been analyzed in the existing literature, including Hart (1995, Chapter 2) in the simplest case without human capital, as well as by Bhattacharyya and Lafontaine (1995) and Blair and Lafontaine (2005). The novelty of our paper arises from considering $M \geq 1$.

If the principal chooses $a^i \in \mathbb{R}_+$, then it incurs cost $F^i\left(a^i\right) = \frac{1}{2}\left(a^i\right)^2$. On the other hand, if the agent chooses a^i, then it incurs the cost $f^i\left(a^i\right) = \theta^i F^i\left(a^i\right)$, where $\theta^i > 0$ for all i. Thus, $\theta^i < 1$ (respectively, $\theta^i > 1$) indicates that the agent (respectively, the principal) has a cost advantage in choosing a^i. For example, the principal may have economies-of-scale advantages over individual agents when incurring the cost associated with some transferable actions (e.g. volume discounts in purchasing equipment) or better information regarding the impact of those transferable actions on revenues due to access to more data (e.g. Uber and Lyft when setting prices for rides). In other contexts, the cost or information advantages lie with the agent (e.g. sharecroppers may have better knowledge than the landowners for determining expenditure on seeds, fertilizer and pesticides; the same may be true for franchisees when choosing staff benefits).

This set-up should be interpreted as a simplified version of a more general model, in which all $M + 2$ actions are transferable in principle. In such a framework, if the principal has a very large cost disadvantage in choosing q (and possibly some of the a^i's), then control over these actions will always be given to the agent. Similarly, if the agent has a very large cost disadvantage in choosing Q (and possibly some of the a^i's), these actions will always be controlled by the principal. The transferable actions $\mathbf{a}$ can then be viewed as actions for which the agent and the principal have comparable (but possibly different) costs.

While we have assumed $(\mathbf{a}, q, Q)$ are all costly actions that increase revenues (e.g. investments in advertising, equipment, technology infrastructure, etc.), as discussed later, our main results remain unchanged when we add any number of costless actions, in which the revenue function R is single-peaked (e.g. price, the choice of a particular design out of several options). Moreover, in an Online Appendix we show that similar results to those in this section can be obtained with a more general specification in which costs and revenues have general functional forms, and there is an arbitrary number of non-transferable actions (i.e. more than one of each type). Provided cost functions increase faster in the various actions than the revenue function, the optimal levels of the actions can still be well defined. So, for instance, we can allow for cost functions to be linear in one or more action, provided revenue is strictly concave in these actions. This may fit some examples in Table 1 better. We chose to focus on linear revenue and quadratic costs for expositional clarity.

A key assumption in our specification is that only the realized revenue R is contractible, whereas the underlying actions $(\mathbf{a}, q, Q)$ are not. Note our results would remain unchanged if we added an arbitrary number of contractible decisions (e.g. price) that impact the revenue function and which the principal could set at the same time it chooses the optimal control allocation and contract for the agent.

We also assume that the principal cannot commit to "throwing away" revenue in case a target specified ex-ante is not reached (Holmstrom, 1982). Ex-ante commitments to destroy revenue seem unrealistic, as they require enforcement by an external third party, who then becomes itself subject to a moral hazard problem. This is one reason why such commitments are seldom used in practice (Eswaran and Kotwal, 1984).

Thus, the principal chooses the set $D \subset \{1, .., M\}$ of transferable decisions over which it keeps control, leaving the agent to control decisions $i \in \{1,..,M\} \setminus D$. It offers a revenue-sharing contract (t,T) to the agent, where T is the fixed fee collected by the principal (which can be positive or negative) and $t \in [0,1]$ is the share of revenue kept by the principal. This means the net payoff received by the principal is $tR+T$ and the net payoff received by the agent is $(1-t)R-T$.[4]

In the Online Appendix, we show that restricting attention to such two-part linear contracts is without loss of generality in our set-up. This is because, in the absence of uncertainty, the only thing that matters is the slope of the contract at the equilibrium values of the choice variables. As a result, the optimal outcome can always be replicated with a linear contract (t,T). This property is general and does not depend in any way on our assumption of linear revenues and quadratic costs. It is an extension of similar results obtained in Romano (1994) and Bhattacharyya and Lafontaine (1995) to the case with any number of transferable actions. Note also that linear contracts are prevalent in all of the examples listed in Table 1 (Bhattacharyya and Lafontaine, 1995, and Lafontaine and Shaw, 2005, provide empirical evidence in the contexts of franchising and sharecropping).

We assume the principal holds all the bargaining power. This implies that it always sets (t,T) so that the agent is indifferent between participation and its outside option, which for convenience we normalize to zero throughout. In general, the optimal contract will have different values of (t,T) depending on the allocation of control over transferable actions D. Thus, it is possible for T to be negative under some allocations (i.e. the agent receives a fixed wage) and positive under other allocations (i.e. the agent pays a fixed fee). In Section 6.1, we will consider the case when the agent has a liquidity constraint, so that the principal cannot extract a fixed fee from the agent.

The game that we study has the following timing. In stage 1, the principal chooses the allocation of control over transferable actions $D \subset \{1,..,M\}$ and the associated contract (t,T); the agent decides whether or not to accept the contract. In stage 2, the principal chooses Q and all a^i's such that $i \in D$, while the agent simultaneously chooses q and all a^i's such that $i \in \{1,..,M\} \setminus D$. Finally, in stage 3, revenues $R(\mathbf{a},q,Q)$ are realized; the principal receives $tR+T$ and the agent receives $(1-t)R-T$.

Since the principal extracts the entire expected surplus, given an allocation of decision rights $D \subset \{1,..,M\}$, the principal's profits can be written as

$$\Pi^*(D) = \max_{t,a^1,..,a^M,q,Q} \left\{ \sum_{i=1}^{M} \beta^i a^i + \phi q + \Phi Q - \sum_{i \in D} \frac{1}{2}\left(a^i\right)^2 - \sum_{i \in \{1,..,M\} \setminus D} \frac{\theta^i}{2}\left(a^i\right)^2 - \frac{1}{2}q^2 - \frac{1}{2}Q^2 \right\}$$

$$\text{s.t.} \quad \begin{cases} a^i = t\beta^i \text{ for } i \in D \\ a^i = \frac{(1-t)\beta^i}{\theta^i} \text{ for } i \in \{1,..,M\} \setminus D \\ q = (1-t)\phi \\ Q = t\Phi. \end{cases}$$

[4] In our model it is immaterial whether the principal or the agent collects revenues R and pays the other party their share according to the contract (t,T). For instance, if the principal is a firm that employs the agent, then the contract can be interpreted as a combination of fixed wage plus bonus in an employment relationship.

In general, for any D, the principal's profits are lower than the first-best profit level

$$\max_{a^1,..,a^M,q,Q} \left\{ \sum_{i=1}^{M} \beta^i a^i + \phi q + \Phi Q - \sum_{i=1}^{M} \frac{\min\left\{\theta^i, 1\right\}}{2} \left(a^i\right)^2 - \frac{1}{2} q^2 - \frac{1}{2} Q^2 \right\}.$$

This reflects two sources of inefficiency. First, the party given control over a^i may not be the more efficient party at carrying it out (that party is the agent if $\theta^i < 1$ and the principal if $\theta^i > 1$). Second, revenue needs to be divided between the principal and the agent to incentivize each of them to choose their respective actions. This inefficiency is the moral hazard in teams identified by Holmstrom (1982), where a team here consists of the agent and the principal.[5] These two inefficiencies create a tradeoff. On the one hand, one would like to allocate transferable actions based on cost advantage; on the other hand, splitting the transferable decisions based solely on the relative cost advantages would ignore the double-sided moral hazard, thus exacerbating the revenue-sharing inefficiency. As a consequence, if cost differences between the principal and the agent are not too large, it may be more efficient to allocate all transferable decisions to the same party, even if this results in some decisions to be mis-allocated from a cost advantage perspective. The following result formalizes this insight.

Proposition 1 *If θ^i is sufficiently close to 1 for all $i \in \{1,..,M\}$, it is optimal to give control over all M transferable actions to the same party. I.e. there exist $\left(\varepsilon^1,..,\varepsilon^M\right) \in \mathbb{R}_+^M$ such that, whenever $\left|\theta^i - 1\right| \leq \varepsilon^i$ for all $i \in \{1,..,M\}$, the optimal allocation of control over transferable actions is either $D^* = \{1,..,M\}$ (the principal controls all transferable decisions) or $D^* = \emptyset$ (the agent controls all transferable decisions).*

Proposition 1 says that strictly interior splits of control rights over costly transferable actions should only occur when there are significant cost differences between principal and agent in undertaking those actions. Otherwise, the principal is better off minimizing the distortion due to revenue-sharing by choosing one of the two extreme allocations of decision rights: $D = \{1,..,M\}$, which we call the $\mathcal{P}$-mode (the principal controls all transferable decisions) and $D = \emptyset$ which we call the $\mathcal{A}$-mode (the agent controls all transferable decisions).

To understand the mechanism underlying this result, suppose the agent and the principal have the same costs of undertaking the transferable actions, i.e. $\theta^i = 1$ for all $i \in \{1,..,M\}$. Then giving control rights over all transferable actions to the party that obtains a higher share of revenues in equilibrium (i.e. aligning low-powered and high-powered incentives) reduces revenue-sharing distortions and thereby raises the principal's profit. To see this, denote by t^* the optimal share of revenues extracted by the principal at the optimal allocation of control rights D^*. If $t^* < 1/2$ and $D^* \neq \emptyset$, then the distortions can be reduced by shifting control over all transferable actions in D^* from the principal to the agent. Indeed, this changes the first-order condition determining action a^i in the second stage from $a^i = t^* \beta^i$ to $a^i = (1 - t^*) \beta^i$ for all $i \in D^*$. The first-order conditions in q and Q stay unchanged.

[5] The only way first-best would be attained in our framework is if all investments were transferable and one party (principal or agent) had no cost disadvantage in choosing any investment. That party should then be given control over all investments and 100% of revenue.

Given $1 - t^* > 1/2 > t^*$, this change leads to an outcome that is closer to the first-best (in which $a^i = \beta^i$) and therefore higher equilibrium profits for the principal. If $t^* > 1/2$ and $D^* \neq \{1, .., M\}$, then, by the same logic, profits can be increased by shifting control over all transferable actions not already in D^* to the principal.[6] Finally, by continuity, this result extends to small cost differences between the principal and the agent, i.e. to the case when all θ^i's are close to 1.

The finding that all transferable decisions should be controlled by the same party when cost asymmetries are small is not driven by positive interaction effects between the various non-contractible actions in revenue, nor by any cost economies of scope across transferable actions.[7] Indeed, we have assumed the revenue function is linear in the various actions and the costs of the transferable actions are independent of one another. If there were economies of scope among them, that would provide an additional reason for giving the same party control (and therefore cost responsibility) for all of these actions.

Even when the cost differences between the principal and the agent are significant, the revenue-sharing disadvantage of splitting decision rights can still lead to all transferable actions being chosen by the same party. To illustrate this point, consider an example with $M = 2$ and $\beta^1 = \beta^2 = \phi = \Phi = 1$. Figure 1 shows the regions in the space $(\theta^1, \theta^2) \in [0, 2] \times [0, 2]$ where each mode dominates: light gray for $\mathcal{A}$-mode, dark gray for the hybrid mode and black for $\mathcal{P}$-mode.

Figure 1: Optimal choice of mode with cost differences

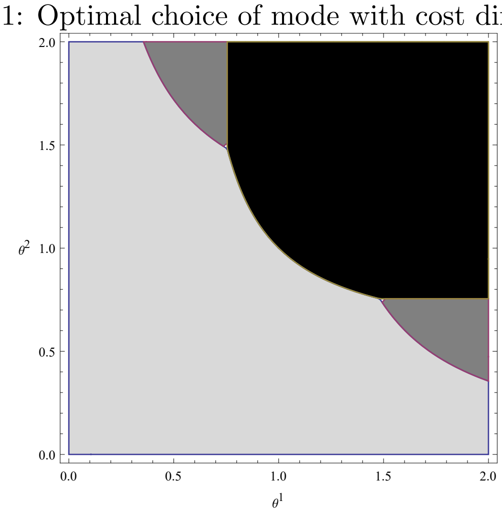

When $\theta^1 < 1$ and $\theta^2 < 1$, the $\mathcal{A}$-mode is obviously optimal. Alternatively, when $\theta^1 > 1$ and $\theta^2 > 1$, the $\mathcal{P}$-mode is obviously optimal. The interesting cases arise when either $\theta^1 < 1 < \theta^2$ or $\theta^2 < 1 < \theta^1$, i.e. where one party is more efficient at choosing a^1 and the other party is more efficient at choosing a^2. As can be seen from Figure 1, the hybrid mode only dominates for a relatively small part of the upper left and bottom right quadrants ($\theta^1 < 1 < \theta^2$ and $\theta^2 < 1 < \theta^1$). Thus, Figure 1 illustrates that

[6] It is straightforward to verify that $t^* = 1/2$ is never possible in this model.

[7] This is why the mechanism underlying our result is different from the one underlying similar results regarding asset ownership in the traditional make vs. buy literature. For example, in Holmstrom and Milgrom (1994), the principal's objective function must be super-modular in the choice variables in order to conclude that asset ownership is correlated with higher incentive payments.

the concern for revenue-sharing distortions over-rides the logic of cost (or information) advantage for a large range of parameter values.

Suppose, in addition to the costly actions, there are some costless transferable actions (e.g. price, horizontal design decisions) over which revenue is concave. Revenue may be higher if set by one party or the other, reflecting that one party may enjoy an information advantage in setting them. Then our analysis above and in what follows in this section will be unaffected. This is because the choice of these costless actions is independent of t, and so it is always efficient to have a^i chosen by the party which has an information advantage in setting it. In particular, there is no reason (based on our model) to expect that they will be allocated to the same party that controls the costly actions.

Let us now focus on the case with symmetric costs for the transferable actions, so $\theta^i = 1$ for $i = 1, ..., M$. From Proposition 1, we know that in this case the optimal allocation of control rights over transferable actions is either $D^* = \emptyset$ ($\mathcal{A}$-mode) or $D^* = \{1, .., M\}$ ($\mathcal{P}$-mode). Let

$$\Pi^{\mathcal{P}*} \equiv \Pi^* (\{1, .., M\}) \text{ and } \Pi^{\mathcal{A}*} \equiv \Pi^* (\emptyset)$$

denote the optimal profits obtained in $\mathcal{P}$-mode and $\mathcal{A}$-mode respectively. We can now derive the following proposition.

Proposition 2 *If $\theta^i = 1$ for all $i \in \{1, .., M\}$, the principal's optimal contract satisfies the following properties:*

1. *If the principal optimally sets $t^* < 1/2$, then the $\mathcal{A}$-mode is strictly optimal (i.e. $\Pi^{\mathcal{A}*} > \Pi^{\mathcal{P}*}$); if $t^* > 1/2$, then the $\mathcal{P}$-mode is strictly optimal (i.e. $\Pi^{\mathcal{P}*} > \Pi^{\mathcal{A}*}$).*

2. *The $\mathcal{A}$-mode dominates the $\mathcal{P}$-mode (i.e. $\Pi^{\mathcal{A}*} > \Pi^{\mathcal{P}*}$) if and only if $\phi > \Phi$.*

The first result in Proposition 2 says that under symmetric costs, the principal would never find it optimal to function in $\mathcal{P}$-mode and keep less than 50% of revenue, or function in $\mathcal{A}$-mode and keep more than 50% of revenue. The second result says that under symmetric costs, the principal prefers the $\mathcal{A}$-mode when revenue depends relatively more on the agent's non-transferable investment than the principal's non-transferable investment, and vice-versa.

Since in the optimal $\mathcal{A}$-mode contract the agent keeps a larger share of variable revenue than in the optimal $\mathcal{P}$-mode contract (i.e. $t^{\mathcal{P}*} > t^{\mathcal{A}*}$), the $\mathcal{A}$-mode is better at generating profit from the agent's non-transferable investments. So when the agent's non-transferable investments are relatively more important in driving revenues (i.e. when $\phi > \Phi$), the principal will prefer the $\mathcal{A}$-mode. The converse is true if the principal's non-transferable investments are relatively more important in driving revenues, i.e. $\Phi > \phi$. By the same reasoning, an increase in ϕ shifts the tradeoff towards the $\mathcal{A}$-mode and an increase in Φ shifts the tradeoff towards the $\mathcal{P}$-mode. Note that the tradeoff between modes does not depend on any of the β^i's, the impact of each transferable action on revenues. The reason is that in both modes the share of revenues retained by the party that chooses the transferable action

a^i ($t^{\mathcal{P}*}$ in $\mathcal{P}$-mode and $\left(1 - t^{\mathcal{A}*}\right)$ in $\mathcal{A}$-mode) is increasing in β^i. Since $t^{\mathcal{P}*}$ and $\left(1 - t^{\mathcal{A}*}\right)$ increase in β^i at the same rate in our linear setting, the resulting tradeoff does not depend on the β.

The first result of Proposition 2 implies that the agent obtains more than 50% of attributable revenues if and only if the principal has chosen (optimally) the $\mathcal{A}$-mode. This provides an empirically testable implication: other things being equal, we expect that organizations that have chosen the $\mathcal{A}$-mode should leave a larger share of their revenues to agents than organizations that have chosen the $\mathcal{P}$-mode.[8] For example, hair salons that rent out chairs charge only a fixed rental fee, letting stylists keep 100% of sales, whereas traditional hair salons that employ their hairstylists offer bonuses ranging from 35% to 60% of sales.[9] Similarly, all sharing economy platforms that rely on independent contractors allow them to keep revenue shares that are significantly above 50% (e.g. between 80% and 90% Postmates and Upwork, 70% for TaskRabbit), whereas their employer counterparts typically only pay their employees fixed wages with more modest bonuses.

If the principal and agent face different costs for transferable actions, Proposition 2 no longer holds for all parameter values. Nevertheless, provided cost differences are not too large, these results are still relevant. To illustrate this, consider an example with $M = 2$ and $\beta^1 = \beta^2 = 1$, and consider the parameter space $\left(\phi, \Phi, \theta^1, \theta^2\right) \in [0, 2]^4$. We find that the $\mathcal{P}$-mode never dominates when agents receive more than 50% of variable revenues (i.e. $t^* < 0.5$), and the $\mathcal{A}$-mode never dominates when agents receive less than 50% of variable revenues (i.e. $t^* > 0.5$), i.e. the first part of Proposition 2 continues to hold. Thus, under these assumptions, a regulator that has to classify workers as either employees ($\mathcal{P}$-mode) or independent contractors ($\mathcal{A}$-mode) and does so on the basis of the observed t^* being smaller or larger than 50%, would never get the classification wrong, even with moderate cost asymmetries.[10]

5 One transferable action and multiple agents

In this section we extend our model to $N > 1$ agents. At the same time, in order to keep things as simple as possible, we assume there is only one transferable action, i.e. $M = 1$, and the principal and the agent face the same cost of carrying out this action.[11] First, we treat the case in which the revenues generated by each agent are independent of one another. Surprisingly, we show that despite this independence, the principal may find it optimal to choose a mixed mode, in which some agents are in $\mathcal{A}$-mode, while others are in $\mathcal{P}$-mode. Next, we allow for spillovers generated by the transferable actions across agents. Focusing on the case in which the principal must choose the same mode for all agents, we study the impact of spillovers on the optimal choice of mode.

[8] In the case of sales agents for products, the percentage commissions for independent agents are higher than those for employees, but both are much lower than 50% of revenues from product sales. This reflects the fact that revenue also includes production costs, which in many cases is not easily observed by the agent. For these cases, the revenue R in our model is best interpreted as revenue net of production costs.

[9] See "Hair & Nail Salons in the US," IBIS World Industry Report 81211, February 2015.

[10] Note that classifying workers as employees or independent contractors when the hybrid mode is optimal cannot be considered a mistake if the regulator is required to classify workers in one of the pure modes.

[11] Recall from Proposition 1 that with symmetric costs, the principal prefers to allocate control rights over a given agent's transferable actions to the same party (the agent or the principal), so $M = 1$ is not a restrictive assumption.

5.1 No spillovers

Let the revenue generated jointly by the principal and agent $i \in \{1, .., N\}$ be

$$R(a_i, q_i, Q) = \beta a_i + \phi q_i + \Phi Q,$$

where a_i denotes the level of the transferable action corresponding to agent i (this choice is made by the agent in $\mathcal{A}$-mode and by the principal in $\mathcal{P}$-mode) and q_i denotes the level of the non-transferable investment chosen by agent i. In this case there is no spillover from one agent's actions to the revenues attributable to other agents. The costs of the various actions are as before: $\frac{1}{2}a_i^2$, $\frac{1}{2}q_i^2$ and $\frac{1}{2}Q^2$ respectively.

We assume that the principal offers each agent i a contract which specifies who (the principal or agent i) controls a_i and a two-part tariff (t_i, T_i). Given that all agents are identical, this means that the principal will effectively choose a pair of two-part tariffs $\left(t^{\mathcal{P}}, T^{\mathcal{P}}\right)$ and $\left(t^{\mathcal{A}}, T^{\mathcal{A}}\right)$, as well as $n \in \{0, .., N\}$, such that all agents $i \in \{1, .., n\}$ receive the $\mathcal{P}$-mode contract $\left(t^{\mathcal{P}}, T^{\mathcal{P}}\right)$ and all agents $i \in \{n+1, .., N\}$ receive the $\mathcal{A}$-mode contract $\left(t^{\mathcal{A}}, T^{\mathcal{A}}\right)$. Following the same logic as in Section 4, restricting attention to such two-part linear contracts is without loss of generality. Thus, the principal controls a_i for all $i \in \{1, .., n\}$, whereas agent i controls a_i for $i \in \{n+1, .., N\}$.

Suppose first that the principal is restricted to pure modes only, i.e. has to choose either $n = N$ (pure $\mathcal{P}$-mode) or $n = 0$ (pure $\mathcal{A}$-mode). In this case, the principal prefers the pure $\mathcal{A}$-mode if and only if $\phi^2 > N\Phi^2$. This is the natural generalization of the second result in Proposition 2 to N agents. It is also a special case of the analysis with spillovers—see (4) with $x = 0$. The factor N in front of Φ^2 reflects that the principal's non-transferable investment Q increases revenue generated jointly with all N agents, whereas agent i's non-transferable investment q_i only increases revenues generated by that agent.

Now suppose the principal can choose any $n \in \{0, .., N\}$. At first glance, given the agents are identical and there are no spillovers across the decisions taken by agents in either mode, one might think the outcome should be the same, namely $n = 0$ if $\phi^2 > N\Phi^2$ and $n = N$ if $\phi^2 < N\Phi^2$. Surprisingly, it turns out that this is not the case: the principal may find it optimal to choose a mixed mode, in which some agents offer their services in $\mathcal{A}$-mode, while the rest offer them in $\mathcal{P}$-mode.

To understand this, note that the fixed fees $T^{\mathcal{P}}$ and $T^{\mathcal{A}}$ are set to extract the entire expected surplus from the agents (as in the case with one agent studied in the previous section). Therefore, the principal's optimal profits conditional on its choice of n can be written

$$\Pi^*(n) = \max_{t^{\mathcal{A}},t^{\mathcal{P}},a_1,..,a_N,q_1,...,q_N,Q} \left\{ \sum_{i=1}^{N} \left(\beta a_i + \phi q_i + \Phi Q - \frac{1}{2}a_i^2 - \frac{1}{2}q_i^2 \right) - \frac{1}{2}Q^2 \right\} \tag{1}$$

s.t.

$$\begin{cases} a_i = t^{\mathcal{P}}\beta \text{ for } i \in \{1,..,n\} \\ a_i = \left(1 - t^{\mathcal{A}}\right)\beta \text{ for } i \in \{n+1,..,N\} \\ q_i = \left(1 - t^{\mathcal{P}}\right)\phi \text{ for } i \in \{1,..,n\} \\ q_i = \left(1 - t^{\mathcal{A}}\right)\phi \text{ for } i \in \{n+1,..,N\} \\ Q = \left(\frac{n}{N}t^{\mathcal{P}} + \frac{N-n}{N}t^{\mathcal{A}}\right)N\Phi. \end{cases}$$

The key observation is that Q is a common investment that increases revenues for *all* agents, e.g. the quality and maintenance of a platform's infrastructure, or advertising conducted by the principal. Thus, the principal's choice of Q is based on $\frac{n}{N}t^{\mathcal{P}} + \frac{N-n}{N}t^{\mathcal{A}}$, the average revenue share collected by the principal from the n agents in $\mathcal{P}$-mode and the $N-n$ agents in $\mathcal{A}$-mode. Fix $t^{\mathcal{P}}$ and $t^{\mathcal{A}}$ such that $t^{\mathcal{P}} > t^{\mathcal{A}}$, which will always be true in equilibrium (indeed, as explained in Section 4, more control should always be associated with a larger revenue share). Since the principal's profit function is concave in Q and the principal's choice of Q is always below the first-best level, this means the principal's objective function is increasing and concave in n. Indeed, substituting one agent in $\mathcal{A}$-mode for an agent in $\mathcal{P}$-mode results in a higher average revenue share obtained by the principal, which increases Q, which in turn leads to higher profits. Due to concavity in Q, the resulting increase in profits is lower when the number of agents in $\mathcal{P}$-mode is larger.

However, there are two other effects associated with shifting some agent i from $\mathcal{A}$-mode to $\mathcal{P}$-mode. The first one is that the agent now invests less in his q_i (the investment is now proportional to $1-t^{\mathcal{P}}$, which is smaller than $1-t^{\mathcal{A}}$ before the switch), which decreases the profits generated by that agent. The second effect is that the investment in a_i shifts from $\beta\left(1-t^{\mathcal{A}}\right)$ to $\beta t^{\mathcal{P}}$, which can be higher or lower. Neither of these two effects depends on the number of agents in $\mathcal{P}$-mode. If the net sum of these two effects on the profits generated by agent i is positive, then clearly the optimal n will be equal to N. On the other hand, if the net sum is negative, then increasing n has a positive effect which is concave in n (through Q) and a negative effect which is linear in n (through q_i and a_i). In the latter case, the optimal n may be interior.

The following proposition confirms this intuition.

Proposition 3 *The optimal number of agents in $\mathcal{P}$-mode is*

$$n^* = \begin{cases} N & \text{if} \quad N\Phi^2 > \beta^2 + \phi^2 \\ N\left(1 - \frac{\phi^2\left(\beta^2+\phi^2-N\Phi^2\right)}{2N\Phi^2\beta^2}\right) & \text{if} \quad \beta^2+\phi^2 \geq N\Phi^2 \geq \phi^2 - \frac{\beta^2\phi^2}{2\beta^2+\phi^2} \\ 0 & \text{if} \quad N\Phi^2 < \phi^2 - \frac{\beta^2\phi^2}{2\beta^2+\phi^2}. \end{cases}.$$

Note that n^* is increasing in $N\Phi^2$ (the importance of the principal's moral hazard) and decreasing in ϕ^2 (the importance of agents' moral hazard), consistent with the intuition built in Section 4. In other words, increasing the importance of the agent's (respectively, the principal's) moral hazard shifts the tradeoff in favor of the $\mathcal{A}$-mode (respectively, $\mathcal{P}$-mode).

Mixed modes, with some agents offering their services in $\mathcal{P}$-mode and others in $\mathcal{A}$-mode, are found quite often in the markets where our theory is relevant. Many consultancies, hair salons and industrial companies relying on sales representatives that use a mix of employees ($\mathcal{P}$-mode) and independent contractors ($\mathcal{A}$-mode). Similarly, many franchisors (e.g. Hertz, InterContinental, McDonalds and Starbucks) use a mix of company-owned and franchised outlets. Furthermore, the empirical studies of franchise chains by Lafontaine and Shaw (2005) and Blair and Lafontaine (2005) show that higher levels of investment by the franchisor (captured by Q in our model) are associated with a larger number of agents in $\mathcal{P}$-mode (higher n^* in our model)—this supports our theoretical prediction above.

The result in Proposition 3 shows that a mixed mode across agents is a strategic way for the principal to get some of the advantages of both pure modes, given that there is some common infrastructure that the principal has to make ongoing investments in (i.e. Q).[12] Such common infrastructure investments are particularly important for online platforms and are in large part responsible (along with network effects) for the extreme scalability of these platforms (see, Levin, 2011). Thus, our model predicts that it is still possible for mixed modes to be optimal even when the principal controls all common infrastructure investments. This is in contrast with Bai and Tao (2000)'s explanation for the optimality of mixed modes, which crucially relies on the existence of a public good problem across agents, i.e. that non-contractible investments by one agent affect the revenues of other agents.

5.2 Spillovers

In the previous sections, the need to share revenues created the distortions that drove our results and tradeoffs. In this section we allow for a second source of distortions: spillovers from the level of transferable actions chosen by one agent on the revenues generated by other agents. Specifically, the revenue function generated jointly by the principal and agent $i \in \{1, .., N\}$ is now

$$R\left(a_i, \overline{a}_{-i}, q_i, Q\right) = \beta a_i + x\left(\overline{a}_{-i} - a_i\right) + \phi q_i + \Phi Q,$$

where $\overline{a}_{-i} \equiv \frac{\sum_{j \neq i} a_j}{N-1}$ is the average of the transferable actions chosen for $j \neq i$. The cost functions remain the same.

We wish to focus on the effect of the spillover x on the principal's optimal choice of mode. To keep the analysis tractable and concise, in this section we assume the principal is restricted to offer the same contract to all agents. This means all agents will be in the same mode ($\mathcal{P}$-mode or $\mathcal{A}$-mode), so that $n = N$ or $n = 0$. This restriction could be justified in cases where having different agents under different organization modes may lead to internal frictions or other costs. Examples of companies using pure modes include the Catalant, TaskRabbit and Upwork marketplaces, and the Dunkin' Donuts,

[12] In contrast, in Hagiu and Wright (2015b) we do not allow for any ongoing investment by the principal, which is why a mixed mode is never optimal in their setting in the absence of spillovers.

RE/MAX and Subway franchises (which are all 100% franchised).

Thus, the principal can choose to either control all transferable actions a_i, $i \in \{1,..,N\}$ (i.e. operate in $\mathcal{P}$-mode) or allow each a_i to be chosen by agent i (i.e. operate in $\mathcal{A}$-mode). In either mode, the principal offers a linear revenue-sharing contract (t, T). The restriction to such two-part linear contracts follows the same logic as before, and is again without loss of generality.

When spillovers are negative ($x < 0$), revenue R is decreasing in $\overline{a}_{-i}$, which means that in $\mathcal{A}$-mode the transferable actions a_i are set too high. Conversely, when spillovers are positive ($x > 0$), revenue R is increasing in $\overline{a}_{-i}$, so that in $\mathcal{A}$-mode the a_i's are set too low. We choose this particular normalization of spillovers, in which x multiplies $\overline{a}_{-i} - a_i$ rather than just $\overline{a}_{-i}$, because it simplifies the analysis.

Consider the following examples from Table 1:

- Hair salons, consulting and outsourcing: spillovers from the marketing of individual professionals are likely negative. A larger investment in the marketing of a given individual professional typically leads to business-stealing from the other professionals.

- Transportation (Uber vs. traditional taxi): spillovers from investments in car quality are likely positive. Better car quality for each individual driver improves the brand image of the entire service in the eyes of customers and therefore helps all other drivers. Business stealing is limited since users rarely, if ever, have the opportunity to choose drivers based on their cars.

- Franchising: spillovers from investments in staff are likely positive. If a given franchisee's staff are more motivated, they provide a better quality of service to customers—this improves the brand image of the franchisor, which in turn helps all other franchisees. Moreover, business stealing among franchisees is limited, since franchisees usually have a certain degree of territorial exclusivity and consumers choose franchisees based on their location.

Note that $R^i \equiv R(a_i, \overline{a}_{-i}, q_i, Q)$ does not depend on the level of non-transferable actions q_j chosen by other agents $j \neq i$, i.e. there are no spillovers resulting from the choices of the agents' non-transferable investments q_j. Allowing such spillovers would not add anything meaningful to the analysis because they would be left uninternalized in either mode (q_j is always chosen by agent j).

The interaction between revenue sharing and spillovers creates the possibility of interesting new results. The revenue-sharing distortion implies that we are in a second-best world in both modes. In this context, positive spillovers lead to the a_i's being set too low in $\mathcal{A}$-mode, which exacerbates the revenue-sharing distortion. On the other hand, negative spillovers lead to the a_i's being set too high in $\mathcal{A}$-mode, which can offset the distortion due to revenue-sharing. As we show formally below, this possibility has counterintuitive implications for the tradeoff between the two modes.

We make the technical assumptions

$$\beta > 0, \quad x < \beta \quad \text{and} \quad x(\beta - x) < N\Phi^2, \tag{2}$$

which ensure that (i) $R(a_i, \overline{a}_{-i}, q_i, Q)$ is increasing in a_i, (ii) all optimization problems are well defined, and (iii) the optimal variable fees in both modes ($t^{\mathcal{P}*}$ and $t^{\mathcal{A}*}$) are strictly between 0 and 1. Note that all $x < 0$ are permissible under (2).

We obtain (all calculations are given in the appendix)

$$t^{\mathcal{P}*} = \frac{\beta^2 + N\Phi^2}{\beta^2 + \phi^2 + N\Phi^2} \text{ and } t^{\mathcal{A}*} = \frac{N\Phi^2 - x(\beta - x)}{(\beta - x)^2 + \phi^2 + N\Phi^2}, \tag{3}$$

and the following proposition.[13]

Proposition 4 *The principal prefers the $\mathcal{A}$-mode to the $\mathcal{P}$-mode if and only if*

$$\left| \frac{\phi^2 x}{\beta} + \beta^2 + N\Phi^2 \right| < \sqrt{\beta^2 (\beta^2 + \phi^2 + N\Phi^2) + \phi^4}. \tag{4}$$

Consider first the baseline case with no spillovers, i.e. $x = 0$. Then the principal prefers the $\mathcal{A}$-mode to the $\mathcal{P}$-mode if and only if $\phi^2 > N\Phi^2$. This result was already noted in (**??**) and corresponds to the standard tradeoff based on giving control to whichever party's moral hazard is more important.

Consider now the tradeoff for general x. If $\beta^2 + N\Phi^2 < \sqrt{\beta^2 (\beta^2 + \phi^2 + N\Phi^2) + \phi^4}$ (which is equivalent to $\phi^2 > N\Phi^2$), so that moral hazard considerations favor the $\mathcal{A}$-mode, then the $\mathcal{A}$-mode is preferred if and only if the magnitude of spillovers $|x|$ is not too large. Indeed, if the magnitude of spillovers is large, the coordination benefits of the $\mathcal{P}$-mode dominate. On the other hand, if $\phi^2 < N\Phi^2$, so that moral hazard considerations favor the $\mathcal{P}$-mode, then the $\mathcal{A}$-mode is still preferred for an intermediate, bounded range of negative spillovers. To understand why, recall that in $\mathcal{A}$-mode, negative spillovers cause the agents to set their a_i's too high relative to what the principal would like them to choose, all else equal. But this implies that in $\mathcal{A}$-mode, negative spillovers help offset to a certain extent the primary revenue distortion, i.e. a_i's being set too low because the party choosing a_i does not receive the full marginal return when $0 < t < 1$. When this offsetting effect is moderately strong (i.e. the magnitude of negative spillovers is not too large), the resulting levels of a_i's are closer to first-best in $\mathcal{A}$-mode than in $\mathcal{P}$-mode, so the $\mathcal{A}$-mode can dominate (this advantage of $\mathcal{A}$-mode must still be traded-off against the moral hazard advantage of the $\mathcal{P}$-mode when $\phi^2 < N\Phi^2$). When the offsetting effect becomes too strong, the resulting levels of a_i's in $\mathcal{A}$-mode are too far above the first-best levels, so the $\mathcal{P}$-mode dominates again.

Inspection of (4) reveals that the range of spillover values x for which the principal prefers the $\mathcal{A}$-mode is skewed towards negative values, consistent with the explanation in the previous paragraph. Positive spillovers cause the a_i's to be set too low in $\mathcal{A}$-mode, which exacerbates the primary revenue distortion. This makes the $\mathcal{A}$-mode relatively less likely to dominate. There still exists a range of positive spillovers for which the $\mathcal{A}$-mode is preferred provided the agents' moral hazard is more important than that of the principal, but that range is smaller than the corresponding range of negative spillovers.

[13] It is straightforward to verify that neither (4), nor the reverse inequality are ruled out by (2). Thus, the proposition identifies a meaningful tradeoff.

The skew towards negative values of x in condition (4) also implies that, if spillovers are moderately negative, then an increase in their magnitude (i.e. a *decrease* in x) shifts the trade-off in favor of the $\mathcal{A}$-mode.[14] This result runs counter to the common intuition, according to which spillovers should always make centralized control (i.e. $\mathcal{P}$-mode in our model) more desirable due to the ability to coordinate decisions. The reason behind this counterintuitive result is that, when spillovers are moderately negative and their magnitude increases, the $\mathcal{A}$-mode levels of a_i's get closer to the first-best level through the offsetting effect described above, so the $\mathcal{A}$-mode becomes relatively more attractive (the $\mathcal{P}$-mode levels of a_i's are unchanged). If spillovers are positive or very negative, then an increase in their magnitude moves the $\mathcal{A}$-mode levels of a_i's away from the first-best level, so the standard effect is restored.

We can interpret this result in the context of one of the examples noted in Section 3, namely consultancies. If promoting an individual consultant steals business from the other consultants in the same consulting firm (negative spillovers), then consultants do too much self-promotion when they are independent contractors ($\mathcal{A}$-mode), relative to what the firm would choose, other things equal. But this effect can help compensate for sub-optimal incentives to invest in marketing whenever the commission paid to consultants is less than 100%. In this context, if the business-stealing effect of self-promotion across consultants is moderate, then an increase in its magnitude can make the $\mathcal{A}$-mode relatively more desirable, by allowing the firm to pay lower commissions while keeping consultants' incentives constant.

Next, we investigate the impact of ϕ^2 and $N\Phi^2$ on the tradeoff between $\mathcal{A}$-mode and $\mathcal{P}$-mode, by considering their effect on the profit differential $\Pi^{\mathcal{A}*} - \Pi^{\mathcal{P}*}$. From (4), this impact seems difficult to ascertain. Fortunately, one can use first-order conditions and the envelope theorem, which lead to simple conditions (see the appendix for calculations).

Proposition 5 *A larger ϕ shifts the tradeoff in favor of $\mathcal{A}$-mode (i.e. $\frac{d(\Pi^{\mathcal{A}*}-\Pi^{\mathcal{P}*})}{d\phi^2} > 0$) if and only if $t^{\mathcal{A}*} < t^{\mathcal{P}*}$. A larger Φ shifts the tradeoff in favor of $\mathcal{P}$-mode (i.e. $\frac{d(\Pi^{\mathcal{A}*}-\Pi^{\mathcal{P}*})}{d(N\Phi^2)} < 0$) if and only if $t^{\mathcal{A}*} < t^{\mathcal{P}*}$.*

Thus, the effects of both types of moral hazard on the tradeoff (and their interpretation) are the same as in the case without spillovers whenever the share of revenues retained by the principal is larger in $\mathcal{P}$-mode, i.e. $t^{\mathcal{P}*} > t^{\mathcal{A}*}$. Namely, increasing the importance of the agent's (respectively, the principal's) moral hazard shifts the tradeoff in favor of the $\mathcal{A}$-mode (respectively, $\mathcal{P}$-mode). The key difference is that now the presence of spillovers makes it possible to have $t^{\mathcal{A}*} > t^{\mathcal{P}*}$ (this was not possible without spillovers). In particular, this arises if and only if

$$\frac{x}{\beta} + \frac{\beta}{\beta - x} < -\frac{\beta^2 + N\Phi^2}{\phi^2}, \tag{5}$$

i.e. if the spillover x is sufficiently negative. Thus, when the inequality in (5) holds, an increase in

[14] Specifically, if $-\beta^2 - N\Phi^2 < x\frac{\phi^2}{\beta} < 0$, then condition (4) is more likely to hold when x decreases.

the importance of the agents' (respectively, the principal's) moral hazard shifts the trade-off in favor of the $\mathcal{P}$-mode (respectively, $\mathcal{A}$-mode).

The interpretation of this counter-intuitive result runs as follows. Negative spillovers partially offset the revenue-sharing distortion in $\mathcal{A}$-mode. As a result, a higher t induces less distortion of the transferable actions a_i in $\mathcal{A}$-mode, so the principal can charge a higher t in $\mathcal{A}$-mode, to the point that $t^{\mathcal{A}^*} > t^{\mathcal{P}^*}$ if spillovers are sufficiently negative. However, when this occurs, agents retain a lower share of revenues in $\mathcal{A}$-mode than in $\mathcal{P}$-mode, so the level of non-transferable effort q_i they choose is *lower* in $\mathcal{A}$-mode. Consequently, when the agents' moral hazard becomes more important in this parameter region, the $\mathcal{P}$-mode becomes relatively more attractive. Similarly, when the principal's moral hazard becomes more important in the same parameter region, the $\mathcal{A}$-mode becomes relatively more attractive.

Finally, the linear example used in this section implicitly assumes that the price to customers is fixed, so is held the same across the two modes, and that there are no production costs. These are not critical assumptions. In an Online Appendix, we show that Proposition 4 remains unchanged even if the principal chooses price along with the fees (t, T) in its contract, and there are production costs. In other words, the trade-off between the two modes remains the same, even though the profit-maximizing price will differ across the two modes (it is higher for the mode generating higher profits).

6 Extensions

In this section we explore several extensions. The first two extend the model with one agent from Section 4. The third one extends the model with multiple agents and spillovers from Section 5.2. In each case, where detailed proofs are needed to establish the results presented, they are provided in the Online Appendix.

6.1 Worker benefits and employees vs. contractors

Our analysis is relevant to current legal and regulatory debates about whether professionals that work through "sharing economy" service platforms (e.g. Handy, Lyft, Postmates, TaskRabbit, Uber) should be classified as employees rather than as independent contractors. All existing legal definitions emphasize the allocation of control rights as the most important factor in determining this issue, which is consistent with our modelling approach. Where to draw the line between employees and independent contractors based on control rights is a legal and notoriously difficult issue, especially since there are usually several relevant control rights involved.

Our model can be used to understand why mis-classifying such service platforms as employers ($\mathcal{P}$-mode) rather than marketplaces relying on independent contractors ($\mathcal{A}$-mode) can be a serious problem. The key cost difference between the two classifications is that a firm classified as an employer must incur the cost of providing certain benefits for its workers (e.g. worker health insurance, worker tax filings, etc.). This cost would otherwise be borne by the workers if the firm were classified as a marketplace. To keep things as simple as possible, suppose that whenever the principal is classified as an employer (correctly if in $\mathcal{P}$-mode and incorrectly if in $\mathcal{A}$-mode), it has to incur cost B to provide the agent with an additional benefit of B. If the principal is classified as a marketplace, the agent can only get the benefit B by incurring the corresponding cost B (or, equivalently, does not get B at all, since the net payoff is zero).

We start by noting that in all variations of the model analyzed in the previous sections, incorporating the benefit B would have no impact on the principal's choice of which party to give control rights to ($\mathcal{A}$-mode vs. $\mathcal{P}$-mode), regardless of who actually incurs the corresponding cost. Indeed, if for example the principal was mis-classified and had to incur the fixed cost B of providing benefits B to the agent in $\mathcal{A}$-mode, then the principal would lower the fixed wage paid to the agent (i.e. $-T$) by B, leaving its profit and the agent's payoff unchanged. Only if the principal had some cost advantage (or disadvantage) in providing the benefits, would the principal's choice of mode be affected.

This irrelevance of B to the choice of mode no longer necessarily holds when the principal's fixed fees must account for a liquidity constraint faced by agents—a very realistic scenario in practice. To illustrate this point, we adapt our model from Section 4 by supposing that the agent cannot be charged an upfront fixed fee because it is liquidity constrained. We capture this by requiring $T \leq 0$. To keep things as simple as possible, we just focus on the case with $M = 1$, so there is a single transferable action. We say that the principal is "correctly classified" if it has to pay the benefit B when it chooses the $\mathcal{P}$-mode (which should be interpreted as employment), but not when it chooses the $\mathcal{A}$-mode (which should be interpreted as independent contracting). In turn, we say that the principal is "mis-classified" if it is required to pay B both in $\mathcal{P}$-mode and in $\mathcal{A}$-mode.

Consider first the case in which the liquidity constraint is not binding when the principal is correctly classified. We illustrate with a numerical example. Suppose that $\beta = 1$, $\Phi = 0.5$, $\phi = 1.5$, $B = 0.5$ and that the outside option gives the agent a payoff of 1.5. Then, since $\phi > \Phi$, the principal prefers the $\mathcal{A}$-mode when it is correctly classified, consistent with the second result in Proposition 2. With these parameter values, the principal would make a loss if it adopted the $\mathcal{P}$-mode. In $\mathcal{A}$-mode it optimally sets $t^{\mathcal{A}} = 0.071$, extracting only a small percentage of the revenue, and it pays the agent a fixed wage of 0.082 (this means $T = -0.082$), so the liquidity constraint is not binding. The principal obtains a profit of 0.134.

Now consider what happens when the principal is misclassified, so it has to pay for B even if it chooses to operate in $\mathcal{A}$-mode. Since the agent no longer incurs the cost of B in $\mathcal{A}$-mode, the principal will be able to increase the share of revenue it extracts and/or reduce the fixed wage it pays to the agent until the agent is once again indifferent between the contract and its outside option. For these parameter values, the principal optimally reduces the fixed wage to zero and increases $t^{\mathcal{A}}$ to 0.233 (without any liquidity constraint it would have kept the optimal $t^{\mathcal{A}}$ equal to 0.071 and instead charged the agent a fixed fee $T > 0$). The principal still prefers the $\mathcal{A}$-mode since the $\mathcal{P}$-mode is not profitable, but its profit is reduced by 34% to 0.088, reflecting that the principal is now forced to recover profits inefficiently through a higher share of variable revenues. There is an equal reduction in total welfare since the agent's surplus remains equal to its outside option regardless of how the principal is classified.

Similar results can arise if the liquidity constraint is binding in $\mathcal{A}$-mode when the principal is correctly classified. To illustrate, suppose now that $\phi = 3$ and the outside option gives the agent a payoff of 4. The other parameters remain unchanged. In this case, the principal still prefers the $\mathcal{A}$-mode, but the liquidity constraint binds, implying $t^{\mathcal{A}} = 0.108$ and $T^{\mathcal{A}} = 0$ (so the agent neither pays a fixed fee nor receives a fixed wage). The principal obtains a profit of 0.967. Now suppose the principal is misclassified. Since the liquidity constraint was already binding in $\mathcal{A}$-mode under the correct classification, it remains binding now that the principal does not have to leave the agent with as much revenue to make it willing to participate (as the agent no longer has to pay for B). The principal responds to the weaker participation constraint by increasing $t^{\mathcal{A}}$ to 0.168. Once again, this pushes the principal further away from the efficient outcome, extracting too much through variable revenues. Its profit falls by 7.1% to 0.898, although this is still higher than the profit it would obtain if it actually changed to $\mathcal{P}$-mode. As before, total welfare falls by an equal amount.

These examples show that when the agent faces a liquidity constraint, mis-classification can cause the principal to recover profits inefficiently through a higher share of variable revenues, resulting in an outcome that is further distorted away from the first-best relative to the outcome under the correct classification. The implication is a significant loss in the principal's profit (as well as in total welfare), which could threaten its viability once other fixed costs of the principal are taken into account. This is a very real concern for many service marketplaces that have emerged in the last few years: some of them have decided to avoid the regulatory risk of being mis-classified by employing the workers who provide services through them ($\mathcal{P}$-mode), and who otherwise would have been independent contractors ($\mathcal{A}$-mode). Examples of companies that have made this decision from their inception include Enjoy, HelloAlfred and Trusted. Others have started with independent contractors and later turned them into employees (e.g. Luxe, Sprig).

6.2 Private benefits

Transferable actions can drive an additional wedge between the two modes when one or both parties derive private benefits from the choice of these actions. Examples of private benefits include the enhancement of individual agents' reputation and outside opportunities by the marketing of their services (e.g. consultants), the improved reputation of the principal, and opportunities to sell additional products or services.

To incorporate such private benefits, we use the model from Section 4 with $M = 1$ and extend it by supposing that the single transferable action a influences some non-contractible outside payoffs, Ba for the principal and ba for the agent, where $b > 0$ and $B > 0$. Extending the second result in Proposition 2 to this case, we obtain that the principal prefers the $\mathcal{A}$-mode to the $\mathcal{P}$-mode if and only if

$$\left((\beta + b)^2 - B^2\right)\phi^2 > \left((\beta + B)^2 - b^2\right)\Phi^2.$$

Thus, the tradeoff captured by the second result in Proposition 2 is robust to the introduction of private benefits: the tradeoff shifts in favor of the $\mathcal{A}$-mode when the agent's moral hazard or private benefit become more important and in favor of the $\mathcal{P}$-mode when the principal's moral hazard or private benefit become more important. In particular, note that $b = B$ implies $\Pi^{\mathcal{A}*} > \Pi^{\mathcal{P}*}$ if and only if $\phi^2 > \Phi^2$; and $\phi^2 = \Phi^2$ implies $\Pi^{\mathcal{A}*} > \Pi^{\mathcal{P}*}$ if and only if $b > B$. Private benefits play a parallel role to moral hazard in determining the choice of the optimal mode.

6.3 Price as the transferable action

In this section, we conduct a similar analysis to the one in Section 5.2, except that now the transferable action is price rather than a costly investment. As pointed out in Section 4, spillovers are necessary in order for a costless action to generate a tradeoff between $\mathcal{A}$-mode and $\mathcal{P}$-mode, so we focus on the case $x \neq 0$ in this section. The revenue generated by agent i is now

$$R\left(p_i, \overline{p}_{-i}, q_i, Q\right) = p_i\left(d + \beta p_i + x\left(\overline{p}_{-i} - p_i\right) + \phi q_i + \Phi Q\right), \tag{6}$$

where $d > 0$ is the demand intercept and $\overline{p}_{-i}$ is the average of the prices chosen for $j \neq i$. The costs of the non-transferable actions remain the same as in Section 4.

To ensure that $R\left(p_i, \overline{p}_{-i}, q_i, Q\right)$ is single-peaked in p_i and that all optimization problems are well defined, we assume

$$2\beta + \max\left\{N\Phi^2, \phi^2\right\} < \min\left\{0, 2x\right\}.$$

Note that these assumptions imply that $\beta < \min\{0, x\}$, as is natural (demand is decreasing in price).

From (6), positive spillovers ($x > 0$) correspond to the usual case with prices: when other agents increase their prices, this increases the demand faced by agent i. Also, one could reinterpret p_i as quantity instead of price, but then the usual case would be captured by negative spillovers ($x < 0$).

Define

$$k \equiv \frac{N\Phi^2\phi^2}{N\Phi^2 + \phi^2} \in (0, |\beta|),$$

which can be viewed as a measure of the combined importance of the agent's and principal's moral hazards (k is symmetric in $N\Phi^2$ and ϕ^2, and increasing in both).

We then obtain the following proposition.

Proposition 6 *The principal prefers the $\mathcal{A}$-mode if and only if*

$$-\frac{4k\left(k + \beta\right)}{k + 2\beta} < x < 0.$$

This result says that the $\mathcal{P}$-mode is preferred if spillovers are positive or very negative. The logic here is somewhat different from the case with costly transferable actions. Given that the transferable action here (price) does not carry any costs, there is no distortion of price in either mode due to revenue-sharing between the principal and each agent. As a result, the variable fee t can be used in both modes to balance double-sided moral hazard (q_i versus Q) equally well. As before, the $\mathcal{P}$-mode has an advantage in internalizing pricing spillovers across the agents' services. However, due to the strategic complementarity between p_i and (q_i, Q), the level of p_i chosen can either offset or compound the effects of double-sided moral hazard, depending on the sign of the spillovers.

When spillovers are negative ($x < 0$), the fact that agents do not internalize spillovers in $\mathcal{A}$-mode can work in favor of the $\mathcal{A}$-mode. Namely, when $x < 0$, the $\mathcal{A}$-mode leads to an excessively high level of p_i, which can help offset the effects of double-sided moral hazard. If this offsetting effect is moderately strong, then the resulting levels of q_i's and Q are closer to first-best in $\mathcal{A}$-mode than in $\mathcal{P}$-mode, so the $\mathcal{A}$-mode dominates. If the offsetting effect is too strong, then negative spillovers over-compensate and the resulting levels of q_i's and Q in $\mathcal{A}$-mode are too far above the first-best levels, so the $\mathcal{P}$-mode is preferred. In contrast, when $x > 0$, the $\mathcal{A}$-mode leads to p_i being set too low, which compounds the effects of double-sided moral hazard. As a result, the $\mathcal{P}$-mode always dominates in that case.

We can interpret the result that negative spillovers across agents' prices can favor the $\mathcal{A}$-mode in the context of the franchising example. When one franchisee increases its price, this may also steer consumers away from other franchisees. As discussed above, this reflects the over-arching importance of the brand of the franchisor, which leads to positive demand externalities. In turn, this implies that independent franchisees tend to set their prices too high relative to what the franchisor would find optimal, so the latter would prefer to control prices (Blair and Lafontaine, 2005, chapter 7). However, our analysis suggests that the excessive prices charged by independent franchisees can help offset the insufficient on-going investments by both the franchisees and the franchisor due to revenue sharing, so giving franchisees discretion over prices will sometimes be preferred.

7 Comparing "control vs. enable" with "make vs. buy"

In this section we explain how the "control vs. enable" strategic decision that we have studied in this paper differs from the classic "make vs. buy" decision, which has been extensively studied in the economics and strategy literatures following Coase (1937) and Williamson (1975). Fundamentally, the key novelty is that the "controlling vs. enabling" decision is about how to allocate control rights over decisions that directly affect customer demand (e.g. pricing, advertising, service) between the focal firm (principal) and its agents (e.g. suppliers, employees). In contrast, in both the "make" and the "buy" modes of organization, the focal firm (the principal in our model) typically maintains full control over all customer-facing decisions and only decides whether to deal with its suppliers through the market (buy) or through vertical integration or hierarchy (make).

As illustrated in Figure 2, "control" can be viewed as equivalent to "make", but "enable" is very different from "buy". In the "enabling" mode of organization, agents interact directly with customers in the sense that agents are given control over customer-facing decisions. Meanwhile, the firm acts as a platform facilitating these interactions (e.g. via its investments in the corresponding infrastructure). The "enabling" mode of organization has become much more prominent in the last decade, mainly due to the rise of online two-sided platforms like Airbnb, eBay, Task Rabbit, Uber, Upwork, etc. This explains why it has only recently started to receive attention in the management and organization literatures.

Figure 2: Comparing different modes of organization

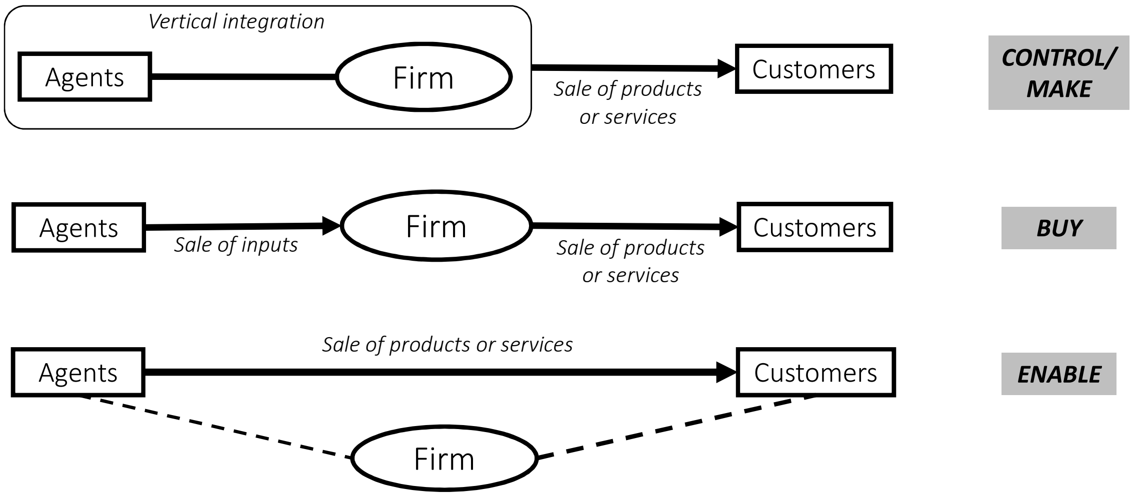

Nevertheless, one may wonder whether the economic and strategic tradeoffs that drive the "make vs. buy" decision are sufficient for also explaining the "control vs. enable" decision? Our analysis and results above clearly show that the answer is no. Let us first start with the common factors, before turning to the differences.

Given our modelling set-up, the most relevant comparison is with the "make vs. buy" (theory of the firm) literature based on property rights (Grossman and Hart, 1986, Hart and Moore, 1990) and incentive systems (Holmstrom and Milgrom, 1994). There are three high-level insights that we share with this literature. First, the prediction that control or ownership should reside with the party whose investments are more important (ex-ante in "make vs. buy", ex-post in "control vs. enable"). Second, the prediction that high-powered incentives (e.g. larger revenue share) should go hand-in-hand with low-powered incentives (control over customer-facing decisions or ownership of assets). Third, the effect of private benefits in our analysis parallels the effect of outside options in Grossman and Hart (1986) and Holmstrom and Milgrom (1994). If one party has larger private benefits or better outside options, that party should be given more control or asset ownership.

Let us now turn to the key differences between "control vs. enable" and "make vs. buy". First, in the "make vs. buy" literature based on property rights and incentive systems, the key instrument determining the choice of organizational mode is the split of asset ownership. This determines the ex-post payoffs earned by the various parties from their respective outside options. Thus, different configurations of asset ownership must lead to different relative configurations of outside options in order for a tradeoff to exist between make and buy. By contrast, in our controlling vs. enabling framework, the key instrument is the allocation of control rights over non-contractible decisions that are chosen ex-post and affect joint payoffs. The tradeoff between the two governance modes is then driven by double-sided moral hazard—outside options are not needed to create a strategic tradeoff (which is why we ignore them in our model).

Second, spillovers created by the transferable decisions corresponding to each agent on the payoffs generated by the other agents are a key feature in "control vs. enable" settings (especially in platform contexts), precisely because the transferable decisions affect customer demand directly. By contrast, the "make vs. buy" literature has largely ignored the role of such spillovers.[15] In the rare instances where such spillovers are mentioned (e.g. Anderson, 1985), they are treated as a form of transaction costs, which leads to the conclusion that they always improve the desirability of the "make" mode relative to the "buy" mode. In contrast, we have seen that (negative) spillovers can make the "enable" mode more desirable relative to the "control" (equivalently, make) mode. Furthermore, spillovers can also lead to reversals of the conventional "make vs. buy" logic that increasing the importance of investments by one party makes it more desirable to give that party control rights.

Third, the key reason why mixed modes between "control" and "enable" can be optimal in our model is the platform nature of the focal firm's investment (i.e. that the principal's investment enhances revenues generated by all agents) combined with revenue sharing. This is entirely novel and very different from the key drivers of mixed modes between "make" and "buy" (Parmigiani, 2007, Puranam et al., 2017): different asset specificities of different goods, uncertainty, scale diseconomies within each mode, and/or cost complementarities between the two modes. None of these factors are present in our model.

Finally, it is worth noting that the notion of "enable" suggested in this paper is a necessary but not a sufficient ingredient for an organization to be a multi-sided platform. Indeed, in addition to enabling direct interactions between two or more groups of users, a multi-sided platform also requires "affiliation" of users on two or more sides to the platform. Affiliation means a specific fixed investment that a user must incur in order to be able to interact with users on the other side(s) (see Hagiu and Wright, 2015b). Thus, except for franchisors that obtain consumer affiliation through loyalty programs, franchisors are not examples of multi-sided platforms even though they fit under the "enable" category in this paper.

8 Conclusions and managerial implications

By substantially reducing the costs of communication and of monitoring revenues generated by independent contractors, Internet and mobile technologies have made it possible to build marketplaces and platforms for a rapidly increasing variety of services. Consequently, the choice facing firms of whether to control the provision of services to customers by employing workers, or whether to enable independent contractors to take control of service provision, and the associated tradeoffs that we have examined in this paper are becoming increasingly relevant in a growing number of industries.

At the most fundamental level, we have shown that the tradeoffs associated with the "control vs. enable" strategy choice arise from the need to balance double-sided moral hazard, while at the same time minimizing distortions in the choice of transferable actions due to revenue sharing. The

[15] Our spillovers are different from Hart and Moore (1990)'s strategic complementaries across investments. Our spillovers affect the tradeoff we study even when an agent's investment decision has no effect on other agents' marginal returns on investment.

first key implication for managers is that low-powered incentives (control over the transferable actions) should be aligned with high-powered incentives (higher share of revenues). In particular, organizations that choose to operate as platforms by giving more control to their agents should also allow these agents to keep a higher share of resulting revenues than organizations that choose to operate in the "control" (vertically integrated) mode. Second, managers should take into account that spillovers across the transferable decisions of different agents introduce an additional distortion. When spillovers are positive (e.g. because one agent's efforts increase the revenues obtained by other agents), the spillover-induced distortion exacerbates the revenue-sharing distortion. Thus, the coordination benefits of "control" shift the baseline tradeoff in favor of the "control" strategy and away from "enable," as standard intuition would suggest. Any increase in such spillovers will make the use of the "control" strategy even more desirable. However, when spillovers are negative (e.g. one agent's efforts decrease the revenues obtained by other agents), they help offset the revenue-sharing distortion. In this case, the tradeoff facing the manager shifts in favor of using an "enable" (platform) strategy and away from "control". As spillovers become even more negative, up to some point, this can further increase the desirability of using the "enable" strategy. Third, when the firm invests in a common infrastructure that helps all of its agents, using a mixed mode across agents can be a strategic way for the manager to get some of the advantages of each of the two pure modes and do better than both.

References

[1] Abhishek, V., K. Jerath and Z.J. Zhang (2015) "Agency Selling or Reselling? Channel Structures in Electronic Retailing," *Management Science*, 62(8), 2259-2280.

[2] Aghion, P. and J. Tirole (1997) "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1-29.

[3] Alonso, R., W. Dessein and N. Matouschek (2008) "When does coordination require centralization?" *American Economic Review*, 98(1), 145–179.

[4] Anderson, E. (1985) "The Salesperson as Outside Agent or Employee: A Transaction Cost Analysis," *Marketing Science*, 4(3), 234-54.

[5] Bai, C. and Z. Tao (2000) "Contract Mixing in Franchising as a Mechanism for Public-Good Provision," *Journal of Economics & Management Strategy*, 9(1), 85-113.

[6] Bester, H. (2009) "Externalities, communication and the allocation of decision rights," *Journal of Economic Theory*, 41, 269-296.

[7] Bhardwaj, P. (2001), "Delegating Pricing Decisions," *Marketing Science*, 20, 143-69.

[8] Bhattacharyya, S. and F. Lafontaine (1995) "Double-Sided Moral Hazard and the Nature of Share Contracts," *RAND Journal of Economics*, 26(4), 761-781.

[9] Blair, R. D. and F. Lafontaine (2005) *The Economics of Franchising*, New York, NY: Cambridge University Press.

[10] Coase, R. (1937) "The Nature of the Firm," *Economica*, 4, 386-405.

[11] Desiraju, R. and S. Moorthy (1997) "Managing a Distribution Channel under Asymmetric Information with Performance Requirements," *Management Science*, 43(12), 1628-1644.

[12] Eswaran, M. and A. Kotwal (1984) "The Moral Hazard of Budget-Breaking," *RAND Journal of Economics*, 15(4), 578-581.

[13] Foros, O., H.J. Kind and G. Shaffer (2013) "Turning the Page on Business Formats for Digital Platforms: Does Apple's Agency Model Soften Competition?" CESifo Working Paper 4362.

[14] Gans, J. (2012) "Mobile Application Pricing," *Information Economics and Policy*, 24, 52-59.

[15] Grossman, S. and O. Hart (1986) "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Ownership," *Journal of Political Economy*, 94(4), 691–719.

[16] Hagiu, A. and J. Wright (2015a) "Marketplace or Reseller?" *Management Science*, 61(1), 184–203.

[17] Hagiu, A. and J. Wright (2015b) "Multi-Sided Platforms," *International Journal of Industrial Organization*, 43, 162–174.

[18] Hart, O. and J. Moore (1990) "Property rights and the nature of the firm," *Journal of Political Economy,* 98, 1119–1158.

[19] Hart, O. (1995) Firms, contracts, and financial structure. Clarendon Press.

[20] Holmstrom, B. (1982) "Moral Hazard in Teams," *Bell Journal of Economics*, 13(2), 324-340.

[21] Holmstrom, B. and P. Milgrom (1994) "The Firm as an Incentive System," *American Economic Review*, 84(4), 972-991.

[22] Jerath, K. and Z.J. Zhang (2010) "Store Within a Store," *Journal of Marketing Research*, 47(4), 748-763.

[23] Johnson, J.P. (2017) "The agency model and MFN clauses," *Review of Economic Studies*, forthcoming.

[24] Lafontaine, F. and K. L. Shaw (2005) "Targeting managerial control: evidence from franchising," *RAND Journal of Economics*, 36(1), 131-150.

[25] Levin, J. D. (2011) "The economics of internet markets," National Bureau of Economic Research Working Paper No. w16852.

[26] Parmigiani, A. (2007) "Why Do Firms Both Make and Buy? An Investigation of Concurrent Sourcing," *Strategic Management Journal*, 28(3), 285-311.

[27] Puranam, P., R. Gulati and S. Bhattacharya (2017) "How Much to Make and How Much to Buy? Explaining Plural Sourcing Strategies," *Strategic Management Journal*, forthcoming.

[28] Romano, R. E. (1994) "Double Moral Hazard and Resale Price Maintenance," *RAND Journal of Economics*, 25(3), 455–466.

[29] Terpstra, A. (1998) *Tenancy and irrigation water management in South-Eastern Punjab, Pakistan*, IIMI Pakistan Report, R-046. Pakistan: International Irrigation Management Institute (IIMI).

[30] Williamson, O.E. (1975) Markets and Hierarchies: Analysis and Antitrust Implications. Free Press: New York.

9 Appendix

9.1 Proof of Proposition 1

For any given (D, t), the principal's profit is equal to

$$\Pi(D,t) \equiv \left(\sum_{i\in D}\left(\beta^i\right)^2 + \Phi^2\right)\frac{t(2-t)}{2} + \left(\sum_{i\in\{1,..,M\}\setminus D}\frac{\left(\beta^i\right)^2}{\theta^i} + \phi^2\right)\frac{(1-t^2)}{2}.$$

Consider first the case with symmetric costs, so $\theta^i = 1$ for $i = 1,..,M$. Then the best profit that the principal can obtain under allocation D is

$$\Pi^*(D) = \max_t\{\Pi(D,t)\} = \frac{\left(\sum_{i=1}^M\left(\beta^i\right)^2 + \Phi^2 + \phi^2\right)^2 - \left(\sum_{i\in D}\left(\beta^i\right)^2 + \Phi^2\right)\left(\sum_{i\in\{1,..,M\}\setminus D}\left(\beta^i\right)^2 + \phi^2\right)}{2\left(\sum_{i=1}^M\left(\beta^i\right)^2 + \Phi^2 + \phi^2\right)}.$$

Since $\sum_{i=1}^M\left(\beta^i\right)^2 + \Phi^2 + \phi^2$ does not depend on D and all the β^i's are positive, it is easily seen that $\Pi^*(D)$ is maximized either by $D^* = \emptyset$ (i.e. the $\mathcal{A}$-mode) or $D^* = \{1,..,M\}$ (i.e. the $\mathcal{P}$-mode). Furthermore, $\Pi^*(D^*) > \Pi^*(D)$ for all D such that $D \neq \emptyset$ and $D \neq \{1,..,M\}$ (note there are $2^M - 2$ such allocations D).

We have proven this result for $\left(\theta^1,...,\theta^M\right) = (1,...,1)$. Because all the functions $\Pi(D,t)$ are continuous in $\left(\theta^1,...,\theta^M\right)$, there exists a neighborhood around $\left(\theta^1,..,\theta^M\right) = (1,..,1)$ such that for any $\left(\theta^1,...,\theta^M\right)$ in this neighborhood, the optimal allocation of control rights remains D^*, i.e. either $\mathcal{A}$-mode or $\mathcal{P}$-mode.

9.2 Proof of Proposition 2

We use the same notation as in the proof of Proposition 1.

Consider first part (1) of the proposition. If $t^* < 1/2$, then (recall $\left(\theta^1,...,\theta^M\right) = (1,...,1)$)

$$\begin{aligned}\Pi(\emptyset,t^*) - \Pi(\{1,..,M\},t^*) &= \left(\sum_{i=1}^M\left(\beta^i\right)^2\right)\frac{\left(\left(1-(t^*)^2\right) - t^*(2-t^*)\right)}{2} \\ &= \left(\sum_{i=1}^M\left(\beta^i\right)^2\right)\frac{1-2t^*}{2} > 0.\end{aligned}$$

Thus, if $t^* < 1/2$, then the $\mathcal{P}$-mode is dominated by the $\mathcal{A}$-mode, which must therefore be optimal (using Proposition 1). By a symmetric argument, if $t^* > 1/2$, then the $\mathcal{A}$-mode is dominated by the $\mathcal{P}$-mode, which is then optimal.

For part (2) of the proposition, straightforward calculations yield

$$\begin{aligned}
\Pi^{\mathcal{P}*} &= \frac{1}{2}\left(\sum_{i=1}^{M}\left(\beta^i\right)^2 + \Phi^2 + \frac{\phi^4}{\sum_{i=1}^{M}\left(\beta^i\right)^2 + \phi^2 + \Phi^2}\right) \\
\Pi^{\mathcal{A}*} &= \frac{1}{2}\left(\sum_{i=1}^{M}\left(\beta^i\right)^2 + \phi^2 + \frac{\Phi^4}{\sum_{i=1}^{M}\left(\beta^i\right)^2 + \phi^2 + \Phi^2}\right).
\end{aligned}$$

Comparing, we have $\Pi^{\mathcal{P}*} > \Pi^{\mathcal{A}*}$ if and only if $\Phi > \phi$.

9.3 Proof of Proposition 3

The principal's problem is defined by (1). Let $\bar{t} \equiv \frac{n}{N} t^{\mathcal{P}} + \frac{N-n}{N} t^{\mathcal{A}}$ denote the "average" transaction fee collected by the principal. After substituting the solutions for a_i, q_i and Q in (1) back into the principal's profit we get

$$\Pi^M\left(t^{\mathcal{P}}, t^{\mathcal{A}}, n\right) = \frac{n\left(t^{\mathcal{P}}\left(2-t^{\mathcal{P}}\right)\beta^2 + \left(1-\left(t^{\mathcal{P}}\right)^2\right)\phi^2\right) + (N-n)\left(\left(1-\left(t^{\mathcal{A}}\right)^2\right)\beta^2 + \left(1-\left(t^{\mathcal{A}}\right)^2\right)\phi^2\right) + \bar{t}\left(2-\bar{t}\right)N^2\Phi^2}{2}.$$

Maximizing this with respect to $\left(t^{\mathcal{P}}, t^{\mathcal{A}}, n\right)$ yields the following first-order conditions (assuming an interior solution in all three variables)

$$\begin{aligned}
&\beta^2 + N\Phi^2 - \left(\beta^2 + \phi^2 + n\Phi^2\right)t^{\mathcal{P}} - (N-n)\,\Phi^2 t^{\mathcal{A}} = 0 \\
&N\Phi^2 - n\Phi^2 t^{\mathcal{P}} - \left(\beta^2 + \phi^2 + (N-n)\,\Phi^2\right)t^{\mathcal{A}} = 0 \\
&\tfrac{\beta^2}{2}\left(t^{\mathcal{P}}\left(2-t^{\mathcal{P}}\right) - 1 + \left(t^{\mathcal{A}}\right)^2\right) + \tfrac{\phi^2}{2}\left(\left(t^{\mathcal{A}}\right)^2 - \left(t^{\mathcal{P}}\right)^2\right) + N\Phi^2\left(1-\bar{t}\right)\left(t^{\mathcal{P}} - t^{\mathcal{A}}\right) = 0.
\end{aligned}$$

Solving the first two first-order conditions above for $\left(t^{\mathcal{P}}, t^{\mathcal{A}}\right)$ as functions of n, we obtain

$$\begin{aligned}
t^{\mathcal{P}} &= \tfrac{\left(\beta^2+N\Phi^2\right)\left(\beta^2+\phi^2\right)+(N-n)\Phi^2\beta^2}{\left(\beta^2+\phi^2\right)\left(\beta^2+\phi^2+N\Phi^2\right)} \\
t^{\mathcal{A}} &= \tfrac{(N-n)\Phi^2\beta^2+N\Phi^2\phi^2}{\left(\beta^2+\phi^2\right)\left(\beta^2+\phi^2+N\Phi^2\right)}.
\end{aligned}$$

This implies

$$\begin{aligned}
t^{\mathcal{P}} - t^{\mathcal{A}} &= \tfrac{\beta^2}{\beta^2+\phi^2} \\
N\left(1-\bar{t}\right) &= \tfrac{(N-n)\beta^2+N\phi^2}{\beta^2+\phi^2+N\Phi^2}.
\end{aligned}$$

We can now plug these expressions into the third first-order condition above, which after simplification becomes

$$-\frac{\phi^2\beta^2}{2\left(\beta^2+\phi^2\right)} + \frac{\Phi^2\beta^2}{\beta^2+\phi^2}\frac{(N-n)\,\beta^2 + N\phi^2}{\beta^2+\phi^2+N\Phi^2} = 0.$$

It is easily verified that the Hessian matrix evaluated at this solution is negative semi-definite so the solution corresponds to a maximum.

Solving for n yields

$$n^* = N\left(1 - \frac{\phi^2\left(\beta^2+\phi^2-N\Phi^2\right)}{2N\Phi^2\beta^2}\right).$$

This solution is valid if and only if

$$0 < \phi^2 \left(\beta^2 + \phi^2 - N\Phi^2\right) < 2N\Phi^2\beta^2,$$

i.e. if and only if

$$\beta^2 + \phi^2 > N\Phi^2 > \phi^2 - \frac{\beta^2\phi^2}{2\beta^2 + \phi^2}.$$

If $N\Phi^2 > \beta^2 + \phi^2$ then $n^* = N$ (pure $\mathcal{P}$-mode is optimal) and if $N\Phi^2 < \phi^2 - \frac{\beta^2\phi^2}{2\beta^2+\phi^2}$ then $n^* = 0$ (pure $\mathcal{A}$-mode is optimal). Note that $\phi^2 - \frac{\beta^2\phi^2}{2\beta^2+\phi^2}$ is increasing in ϕ^2.

9.4 Proof of Proposition 4

Consider first the $\mathcal{P}$-mode. The payoff to agent i is

$$(1-t)\,R_i - \frac{1}{2}q_i^2 - T = (1-t)\left(\beta a_i + x\left(\overline{a}_{-i} - a_i\right) + \phi q_i + \Phi Q\right) - \frac{1}{2}q_i^2 - T,$$

which implies that the level of investment chosen by each agent in the second stage is

$$q^{\mathcal{P}}(t) = \phi\,(1-t)\,.$$

In $\mathcal{P}$-mode, the principal sets $a_1, ..., a_N$ and Q to maximize its second stage revenues (t is set in the first stage):

$$\sum_{i=1}^{N}\left(t\left(\beta a_i + x\left(\overline{a}_{-i} - a_i\right) + \phi q_i + \Phi Q\right) - \frac{1}{2}a_i^2\right) - \frac{1}{2}Q^2,$$

implying the principal's optimal choices are

$$\begin{aligned} a^{\mathcal{P}}(t) &= \beta t \\ Q^{\mathcal{P}}(t) &= N\Phi t. \end{aligned}$$

The fixed fee T is set to render each agent indifferent between working for the principal and her outside option, so the expression of $\mathcal{P}$-mode profits as a function of t is

$$\frac{N}{2}\left(\left(\beta^2 + N\Phi^2\right)t\,(2-t) + \phi^2\left(1-t^2\right)\right). \tag{7}$$

Maximizing (7) with respect to t gives the expression for $t^{\mathcal{P}*}$ in (3) which is positive but smaller than 1. With this optimal fee, the resulting profits in $\mathcal{P}$-mode are

$$\Pi^{\mathcal{P}*} = \frac{N}{2}\left(\beta^2 + N\Phi^2 + \frac{\phi^4}{\beta^2 + \phi^2 + N\Phi^2}\right). \tag{8}$$

Consider next the $\mathcal{A}$-mode. The payoff to an individual agent joining the principal is

$$(1-t)\left(\beta a_i + x\left(\overline{a}_{-i} - a_i\right) + \phi q_i + \Phi Q\right) - \frac{1}{2}a_i^2 - \frac{1}{2}q_i^2 - T.$$

Individual agents maximize their second stage payoff by choosing

$$q^{\mathcal{A}}(t) = \phi(1-t)$$
$$a^{\mathcal{A}}(t) = (1-t)(\beta - x).$$

The principal's second stage profits in $\mathcal{A}$-mode are

$$\sum_{i=1}^{N} t\left(\beta a_i + x\left(\overline{a}_{-i} - a_i\right) + \phi q_i + \Phi Q\right) - \frac{1}{2}Q^2,$$

which the principal maximizes over Q, leading to

$$Q^{\mathcal{A}}(t) = N\Phi t.$$

Stepping back to the first stage, the principal sets T to equalize the agents' net payoff to their outside option. Total profit for the principal in $\mathcal{A}$-mode as a function of t is then

$$\frac{N}{2}\left((\beta - x)(1-t)(\beta + x + (\beta - x)t) + \phi^2\left(1-t^2\right) + N\Phi^2 t(2-t)\right). \tag{9}$$

The optimal variable fee is given by the expression for $t^{\mathcal{A}*}$ in (3). Resulting profits in $\mathcal{A}$-mode are

$$\Pi^{\mathcal{A}*} = \frac{N}{2}\left(\beta^2 - x^2 + \phi^2 + \frac{\left(N\Phi^2 - x(\beta - x)\right)^2}{(\beta - x)^2 + \phi^2 + N\Phi^2}\right). \tag{10}$$

Comparing (8) with (10), the $\mathcal{A}$-mode is preferred if and only if

$$\phi^2 + \frac{\left(N\Phi^2 - x(\beta - x)\right)^2}{(\beta - x)^2 + \phi^2 + N\Phi^2} > N\Phi^2 + x^2 + \frac{\phi^4}{\beta^2 + \phi^2 + N\Phi^2}.$$

If there are no spillovers, i.e. $x = 0$, then this condition simplifies to $\phi^2 > N\Phi^2$. For $x \neq 0$, the condition can be re-written as in (4).

9.5 Proof of Proposition 5

To determine the effects of ϕ^2 and $N\Phi^2$ on the tradeoff between the two modes, we apply the envelope theorem to expressions (7) and (9), and obtain

$$\frac{d\Pi^{\mathcal{P}*}}{d\phi^2} = \frac{N}{2}\left(1 - \left(t^{\mathcal{P}*}\right)^2\right) \text{ and } \frac{d\Pi^{\mathcal{P}*}}{d\left(N\Phi^2\right)} = \frac{N}{2}t^{\mathcal{P}*}\left(2 - t^{\mathcal{P}*}\right)$$
$$\frac{d\Pi^{\mathcal{A}*}}{d\phi^2} = \frac{N}{2}\left(1 - \left(t^{\mathcal{A}*}\right)^2\right) \text{ and } \frac{d\Pi^{\mathcal{A}*}}{d\left(N\Phi^2\right)} = \frac{N}{2}t^{\mathcal{A}*}\left(2 - t^{\mathcal{A}*}\right).$$

Since $0 < t^{\mathcal{P}*}, t^{\mathcal{A}*} < 1$ and $t(2-t)$ is increasing in t for $t \in [0,1]$, we conclude that

$$\frac{d\left(\Pi^{\mathcal{A}*} - \Pi^{\mathcal{P}*}\right)}{d\left(\phi^2\right)} > 0 \text{ if and only if } t^{\mathcal{P}*} > t^{\mathcal{A}*}$$
$$\frac{d\left(\Pi^{\mathcal{P}*} - \Pi^{\mathcal{A}*}\right)}{d\left(N\Phi^2\right)} > 0 \text{ if and only if } t^{\mathcal{P}*} > t^{\mathcal{A}*}.$$