

de Véricourt, Francis; Gurkan, Huseyin

Working Paper

Contracting, pricing, and data collection under the AI flywheel effect

ESMT Working Paper, No. 20-01

Provided in Cooperation with:

ESMT European School of Management and Technology, Berlin

Suggested Citation: de Véricourt, Francis; Gurkan, Huseyin (2020) : Contracting, pricing, and data collection under the AI flywheel effect, ESMT Working Paper, No. 20-01, European School of Management and Technology (ESMT), Berlin

This Version is available at:

<https://hdl.handle.net/10419/215694>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

March 23, 2020

ESMT Working Paper 20-01

Contracting, pricing, and data collection under the AI flywheel effect

Francis de Véricourt, ESMT Berlin

Huseyin Gurkan, ESMT Berlin

Copyright 2020 by ESMT European School of Management and Technology GmbH, Berlin, Germany, <https://esmt.berlin/>.

All rights reserved. This document may be distributed for free – electronically or in print – in the same formats as it is available on the website of the ESMT (www.esmt.org) for non-commercial purposes. It is not allowed to produce any derivatives of it without the written permission of ESMT.

Find more ESMT working papers at [ESMT faculty publications](#), [SSRN](#), [RePEc](#), and [EconStor](#).

Contracting, Pricing and Data Collection under the AI Flywheel Effect

Francis de Véricourt

ESMT Berlin, Germany, francis.devericourt@esmt.org,

Huseyin Gurkan

ESMT Berlin, Germany, huseyin.gurkan@esmt.org,

This paper explores how firms that lack expertise in machine learning (ML) can leverage the so-called AI Flywheel effect. This effect designates a virtuous cycle by which, as an ML product is adopted and new user data are fed back to the algorithm, the product improves, enabling further adoptions. However, managing this feedback loop is difficult, especially when the algorithm is contracted out. Indeed, the additional data that the AI Flywheel effect generates may change the provider's incentives to improve the algorithm over time.

We formalize this problem in a simple two-period moral hazard framework that captures the main dynamics between machine learning, data acquisition, pricing and contracting. We find that the firm's decisions crucially depend on how the amount of data on which the machine is trained interacts with the provider's effort. If this effort has a more (resp. less) significant impact on accuracy for larger volumes of data, the firm underprices (resp. overprices) the product. Further, the firm's starting dataset, as well as the data volume that its product collects per user, significantly affect its pricing and data collection strategies. The firm leverages the virtuous cycle less for larger starting datasets and sometimes more for larger data volumes per user. Interestingly, the presence of incentive issues can induce the firm to leverage the effect less when its product collects more data per user.

Key words: Data, Machine Learning, Pricing, Incentives and Contracting

History: March 3, 2020

1. Introduction

To train ML algorithms, companies often deploy their artificial intelligence (AI)-based products early and collect usage data from their first customers. As new data are fed back to the algorithm, the technology improves, enabling further adoptions and thus the acquisition of additional data. This virtuous feedback loop, sometimes referred to as the

AI Flywheel effect in the popular press (Trautman 2018), compounds the economic effect by which quality increases demand, according to the statistics principle whereby data improves accuracy.

The AI Flywheel effect has many applications, from voice recognition systems (Sarikaya 2019) to self-driving vehicles (Miller 2016), and even explains how certain web search engines ended up dominating the market. Nonetheless, this virtuous cycle is perhaps most useful for smaller teams or novel and specialized applications for which data is scarce. The founders of startup Blue River Technology famously established their first dataset manually to train an AI system that would distinguish weeds from crops, a crucial step for the optimal spray of pesticides in farming (Ng 2018, Trautman 2018). This yielded an algorithm with low performance, but with its adoption by early users, the company could leverage the effects and significantly improve the algorithm. The company was sold in 2017 for more than \$300 million (Golden 2017).

However, despite its apparent simplicity, the AI Flywheel effect is difficult to implement, especially among the small organizations that would most benefit from it. First, a firm makes choices, and pricing decisions in particular, that affect demand alongside accuracy and hence interfere with the virtuous cycle. More specifically, the AI Flywheel effect introduces an additional tradeoff between improving algorithms' accuracy and maximizing revenue, which the firm needs to consider when setting its pricing strategy.

Second, and perhaps more importantly, many firms lack the expertise to develop ML algorithms. Indeed, the economy has experienced a significant shortage of skilled data scientists, which particularly affects start-ups and small organizations (Nicolaus Henke et al. 2016). This shortage has given rise to a striving outsourcing industry (Research Nester Pvt. Ltd 2019), and many start-ups have been successful by turning to technology outsourcing (examples include Skype, Opera and Slack, to name a few; see Cengiz 2015).

However, relying on outsourcing gives rise to incentive issues, which may impair accuracy and thus again interfere with the AI Flywheel effect. For example, a provider may shirk by applying standard third-party software that may be suboptimal for the task or expose the firm to threats (Kendra et al. 2019, Bursztein 2018). If these algorithms are not developed with care, they learn only surface statistical regularities, which affects their ability to generalize and thus their accuracy (Jo and Bengio 2017). More generally, the provider may have an incentive to wait for more usage data before exerting any effort to improve the

algorithm. In fact, incentive issues such as these may not disappear when the firm does not outsource the algorithm but instead employs an expert. Indeed, the details of ML algorithms and their outputs suffer notoriously from a lack of explainability (Lipton 2016, Ribeiro et al. 2016), rendering the expert’s efforts to improve accuracy difficult to observe and contract on.

Third, the amount of data to which the provider has access for training the algorithm may exacerbate the incentive issues. For instance, research in AI has suggested that accuracy depends less on the specifics of the algorithm and more on the data on which it is trained as data volume increases (Banko and Brill 2001, Halevy et al. 2009). In this case, the provider’s effort to develop the machine matters more when data are scarce, i.e., in the early stages of the AI Flywheel effect. Hence, the intensity of the incentive issues may change over time as the feedback loop between accuracy and usage data unfolds.

The goal of this paper is to shed light on how firms that lack expertise in ML can leverage and optimize the AI Flywheel effect. In this context, we seek to understand how the need to mitigate the incentive issues created by outsourcing the algorithm affect the firm’s pricing and data collection strategies. To that end, we formalize the problem in a simple two-period moral hazard framework, which captures the previous three features: the accuracy vs. the revenue tradeoff, the incentive issues, and the impact of data on the intensity of these issues.

Our analysis reveals that the firm’s decisions crucially depend on how the amount of data on which the machine is trained interacts with the provider’s effort. Specifically, if the provider’s effort has a more significant impact on accuracy for larger volumes of data, the firm *underprices* the product in order to acquire more data from the market, retrain the algorithm and generate more revenues in the future. In contrast, the firm *overprices* and collects less data if the provider’s effort is most impactful when data are scarce. In addition, when effort is most impactful for medium amounts of data, the firm’s pricing strategy depends on the initial dataset on which the algorithm is first trained. The firm underprices if the initial dataset is small and overprices otherwise.

This last point reveals the importance of the dataset on which the algorithm is first trained. In particular, firms with different initial datasets may follow radically different pricing strategies. In addition, we provide sufficient conditions under which increasing the

size of this initial dataset reduces the inefficiencies due to incentive issues when the firm underprices but exacerbates them when the firm overprices.

Our analysis also reveals the key role that the data volume per user has on the firm's decisions. Indeed, the firm can sometimes increase the amount of data its product collects per user. For instance, the firm may rely on third-party services in the case of mobile and web applications (Deshpande 2019) or increase the capacity of embedded sensors in case of physical products (McGrath and Scanail 2013). This, in turn, should provide the firm with more data overall to retrain and improve the algorithm. We find, however, that this intuition does not always hold when the firm needs to manage the provider's incentives. If the provider's effort is most impactful when data are scarce, the effect is actually reversed and the firm collects overall *less* data when the data volume per user is sufficiently high.

Taken together, these results characterize how the starting dataset as well as the amount of usage data collected by the product significantly affect the firm's pricing and data collection strategies. The key driver for this lies in the type of impact that data has on the provider's effort when training the algorithm.

More specifically, we consider a firm (the principal) that outsources the development of an ML algorithm to a provider (the agent). (See Section 2.) The expected accuracy of the algorithm increases in both the amount of data on which it is trained and the effort of the provider. At the beginning of the time horizon, the firm starts with a small initial dataset, which allows the provider to develop a first version of the algorithm. The firm chooses a pricing strategy and begins to market this product. Demand for the product decreases with price and increases with accuracy. At the end of this first period, the realized demand generates revenues for the firm and additional data for the provider. This additional amount of data corresponds to the realized demand multiplied by the volume rate, which is the expected amount of data per user that the product collects. The provider retrains the algorithm based on the augmented dataset, and the firm markets this second version at a new price. However, exerting effort to improve the algorithm is costly. Because effort is non-contractable, the firm faces a moral hazard problem in each period when the provider trains the algorithm. The firm then seeks to design a contract and a pricing strategy that maximize its total profit over the time horizon.

A crucial feature of our setup is how effort and data interact with one another to determine the algorithm's accuracy. We introduce the notion of data impact (in Section 2.2),

which maps the size of the available dataset to the normalized effect of shirking on the probability of high accuracy. The data impact characterizes how data interact with effort to determine the algorithm's accuracy. When the data impact is constant in the data size, effort and data independently affect accuracy. However, when the data impact increases (resp. decreases), exerting effort increases accuracy relatively more (resp. less) with more data.

Overall, this setup captures the elementary dynamics of the AI Flywheel effect. Indeed, demand in the first period generates additional data that improve the algorithm's accuracy in the second period, which in turn generates more demand. The model also captures how the firm's decisions affect the flywheel effect in a simple manner. Specifically, the firm faces a price/data tradeoff in which decreasing the price generates more data but less revenues. Further, the initial data size and the volume rate characterize together the potential strength of the AI Flywheel effect. The initial data size specifies the firm's starting point in the virtuous cycle, while the volume rate influences the speed at which this cycle unfolds. Finally, the data impact captures how the intensity of the moral hazard problem changes with data and thus as the AI Flywheel effect unfolds.

In this setup, we begin by solving the first-best problem, that is, the existence of the AI Flywheel effect but without incentive issues (see Section 3). In this case, we find that the flywheel effect induces the firm to decrease its price. We then characterize the optimal contract and equilibrium payments to the provider in the presence of moral hazard. Under this optimal contract, the firm always prices at the first-best price when the data impact is constant in the size of the dataset, but always prices below (resp. above) the first-best price to collect more (resp. less) data when the data impact increases (resp. decreases). (See Section 5.) This situation implies that when the data impact is unimodal, a threshold exists for the size of the starting dataset such that the firm prices below (resp. above) first-best when the starting dataset size is below (resp. above) the threshold. Further, the firm's price always increases in the initial data size. In this case, and under additional mild conditions, the firm underprices less as the initial data increase if the data impact also increases but overprices more when the data impact decreases. (See Section 6.) Finally, we show that when the data impact is increases, improving the volume rate also increases the amount of collected data. In contrast, improving the volume rate has a unimodal effect on the collected data when the data impact decreases (see Section 7).

1.1. Literature Review

The advent of the digital economy has recently generated new research on data privacy and markets in both management science and economics. This new stream of research explores the impact of data leakage in platform business models (Acemoglu et al. 2019), the issue of selling data (Bimpikis et al. 2019, Mehta et al. 2019), and the effect of collecting data on privacy and price discrimination (Loertscher and Marx 2019). In contrast, our work focuses on the outsourcing of ML algorithms that make use of this data, which creates incentive issues that dynamically interact with the amount of available data. More generally, ours is the first study in this stream of research to explore the problem of managing the AI Flywheel effect.

Our work is also related to the large literature on dynamic pricing with learning. Of particular interest is the recent stream of research on new experience goods and quality learning. Yu et al. (2015), for instance, study the dynamic pricing of new experience goods in the presence of two-sided learning (learning about quality via consumer reviews). In their setting, the pricing decision affects both revenue and the flow of information. They show that consumer-generated quality information may decrease the firm's profit and even consumer surplus. Feldman et al. (2018) also analyze the pricing and quality design of new experience goods for consumers who are social learners. They characterize the deviation of the firm's optimal policy from a setting without social learning. In these setups, quality is a decision variable that is set *ex ante*, and learning concerns either the firm learning about the consumers or the consumers learning about quality. By contrast, accuracy is dynamically improved in our setting and learning concerns the training of the algorithm, *i.e.*, the enhancement of quality. In addition, we consider moral hazard issues, which is not the focus of this stream of research.

From a more technical point of view, our model is a dynamic moral hazard problem with binary effort choices and binary outcomes. Different versions of this problem have been studied, especially in the sales force management literature. Schöttner (2016) analyzes a multi-period setting with different sales probabilities in each period when the firm can obtain only aggregate information on sales. Despite being different across periods, the sales probabilities are taken as constant. Kräkel and Schöttner (2016) consider a two-period model with binary effort choices and analyze the optimal contracts. They also study a case where the second-period sales opportunity randomly depends on the outcome of the first

period with exogenous probabilities. Schmitz (2005) explore a similar setting, where the probability of a favorable outcome in the second period takes a larger value in the case of a success in the first period.

Although these papers explore various configurations of the uncertainty structure, the probability of observing a favorable outcome in their models is fixed (Schöttner 2016) or exogenously depends on the outcome in the first period (Kräkel and Schöttner 2016, Schmitz 2005). By contrast, the outcomes in the first period endogenously determine the probability of success in the second period in our setup. This feature is indeed at the core of the AI Flywheel effect.

In this stream of research, Dai and Jerath (2019) study a slightly more general dynamic moral hazard problem with binary effort choices and three possible outcomes. A key aspect of this model, also introduced in de Véricourt and Gromb (2018, 2019) for more general distributions of outcomes, is that the firm’s capacity decision interacts with the moral hazard problem. Indeed, in de Véricourt and Gromb (2018), Dai and Jerath (2019) and de Véricourt and Gromb (2019), a low capacity level may censor high demand realizations, which exacerbates the incentive issue. Our work considers a different type of interaction between the firm’s decision (pricing in our case) and the moral hazard problem through a monotonic property of the data impact ratio, which is more appropriate in the context of training an algorithm with data.

Finally, our work also contributes to the rich operations management literature in entrepreneurship, as our setup is particularly relevant for cash-constrained firms with a lack of technical skills. The points of focus in this literature widely range from investment timing (Swinney et al. 2011) and financial capabilities (Tanrisever et al. 2012) to complementary technologies (Anderson Jr and Parker 2013). By contrast, we provide insight on how cash-constrained firms can leverage a business model based on the AI Flywheel effect.

2. Model Description

We model the problem of managing the AI Flywheel effect in an elementary two-period moral hazard framework. In our setup, the firm (the principal) outsources the development and training of the algorithm to a provider (the agent) in the beginning of each period. The resulting accuracy of the algorithm depends on both the provider’s effort and the size of the available dataset. Given this accuracy level, the firm markets the product to users with heterogeneous accuracy and price sensitivities. At the end of each period, demand is realized, which determines the profit and the additional generated data for that period.

2.1. Data, Accuracy and Revenue

At the beginning of period $t = 1, 2$, the size of the available dataset to train the algorithm is equal to d_{t-1} . In particular, d_0 denotes the size of the firm's initial dataset. In each period t , the provider can either exert effort $e = w$ at cost $\kappa > 0$ or shirk $e = s$ at no cost. We denote by α_t , the resulting algorithm's accuracy, which can be either high or low with $\alpha_t \in \{\alpha_h, \alpha_\ell\}$, where $\alpha_h > \alpha_\ell$. Given data size d_{t-1} and effort e , the probability of high accuracy ($\alpha_t = \alpha_h$) is equal to $\pi_e(d_{t-1})$ with $\pi_w(d_{t-1}) > \pi_s(d_{t-1})$. Probabilities $\pi_e : [d_{\min}, d_{\max}] \rightarrow [0, 1)$, $e \in \{s, w\}$, are twice-differentiable and increasing functions with a continuous second-order derivative. Moreover, π_w is concave. In other words, the accuracy of ML models increases with data, but the marginal effect of additional data is decreasing (see Banko and Brill 2001, for instance). The further lower bound $d_{\min} > 0$ is the minimum size required for the development of a functioning ML model, and $d_{\max} < \infty$ is the largest possible total data size on which the algorithm can be trained.

Given accuracy α_t , the firm then prices and markets the product. The market in period t corresponds to a continuum of buyers of total mass normalized to one, a common framework in the pricing literature (e.g., Aflaki et al. 2019, Feldman et al. 2018, Yu et al. 2015). Each buyer has a private accuracy sensitivity v that is drawn from the standard uniform distribution with c.d.f. F , p.d.f. f , support $[0, 1]$, and a virtual value function $\phi(v) = v - \bar{F}(v)/f(v)$, where $\bar{F}(\cdot) = 1 - F(\cdot)$. Hereafter, we use the notation \bar{x} to denote $1 - x$ for an arbitrary term x . A buyer with sensitivity v purchases the product with accuracy α and price p if $\alpha v - p \geq 0$. Given accuracy-price pair (p, α) , demand is equal to $\bar{F}(p/\alpha)$, which yields revenue $p\bar{F}(p/\alpha)$. Equivalently, the firm may choose demand quantity q instead of price p , with $p = \alpha F^{-1}(\bar{q})$ and, hence, $p\bar{F}(p/\alpha) = q\alpha F^{-1}(\bar{q})$, where F^{-1} is the inverse of the c.d.f. F .

One key feature of our model is that demand generates data that can be used to improve the algorithm. We denote by volume rate $\nu > 0$ the average amount of usage data that the product can collect, i.e., the marginal data generated per user. Given demand q , therefore, the total amount of collected data δ in the period is equal to $\delta = \nu q$. Thus, at the end of period $t = 1, 2$ the size of the available dataset is equal to $d_t = \delta_t + d_{t-1}$. Further, the firm's revenue can be expressed in terms of volume rate ν , collected data δ and accuracy α . Specifically, we define this revenue as $R_\nu(\delta, \alpha)$ with $R_\nu(\delta, \alpha) \triangleq \alpha\delta/\nu F^{-1}((1 - \delta/\nu))$.

Thus, decreasing price increases demand and hence the amount of collected data but may also decrease revenue R_ν . In this sense, our model captures the tradeoff associated with the AI Flywheel effect between maximizing revenue and collecting additional data. In addition, parameters (d_0, ν) characterize the potential strength of the AI Flywheel effect in our setup. Indeed, we have $d_1 = \nu q + d_0$; hence, data size d_0 specifies the firm's starting point in the virtuous cycle, while volume rate ν influences the speed at which the firm can leverage this cycle.

Figure 1 depicts the timing of the events corresponding to our setup.

Period 1a - Initial Algorithm Development. The firm starts with an initial dataset of size d_0 . Based on this dataset, the provider chooses effort $e \in \{s, w\}$ to develop a first version of an ML algorithm. The algorithm's accuracy α_1 is then realized according to probability $\pi_e(d_0)$.

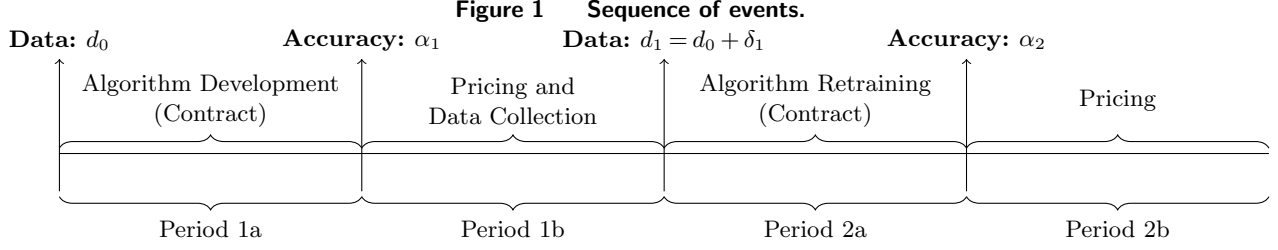
Period 1b - Pricing and Data Collection. Given accuracy α_1 , the firm prices and markets the product. The firm collects δ_1 such that the total size of the dataset becomes $d_1 = d_0 + \delta_1$ and generates revenue $R_\nu(\delta_1, \alpha_1)$.

Period 2a - Algorithm retraining and improvement. If $\alpha_1 = \alpha_\ell$, the provider retrains the algorithm with an augmented dataset of size d_1 . (Otherwise, the maximum possible accuracy level α_h is achieved and the firm does not need the provider to improve accuracy further.)¹ The provider again chooses effort e to retrain the algorithm, which yields accuracy α_2 according to probability $\pi_e(d_1)$. The probability of achieving high accuracy α_h increases in this period, i.e., $\pi_e(d_1) > \pi_e(d_0)$, because of the dataset increase $d_1 \geq d_0$, and since more data improves accuracy, $\pi_e(\cdot)$ for $e \in \{s, w\}$ are increasing functions.

Period 2b - Pricing. Finally, given accuracy α_2 , the firm prices and markets the product, which determines δ_2 and generates revenue $R_\nu(\delta_2, \alpha_2)$.

At the end of the time horizon, the firm has no further incentive to retrain the algorithm using additional data δ_2 , which also means that d_1 corresponds to the size of the largest dataset on which the algorithm is ultimately trained. The maximum possible size of this dataset is then equal to $d_{\max} = \nu + d_0$ (recall that the market size is normalized to one). In the following, we thus take $d_{\min} \leq d_0 \leq d_{\max}/2$ and $0 < \nu \leq d_{\max}/2$.

¹ Our model and results easily extend to the case where accuracy is cumulative, that is, in situations where accuracy can be further improved when $\alpha_1 = \alpha_h$.



2.2. The Data Impact

Thus far, we have ignored the incentive issues that outsourcing the development of the algorithm creates. In particular, the accuracy level depends on the available data and the provider's effort. Thus, the availability of data may interact with the intensity of the moral hazard problem that the firm faces in each period. To characterize this interaction between effort and data, we introduce the notion of the *data impact*, which we denote by $\rho(d)$. The data impact maps dataset size d to the normalized effect of shirking on the probability of high accuracy, i.e.,

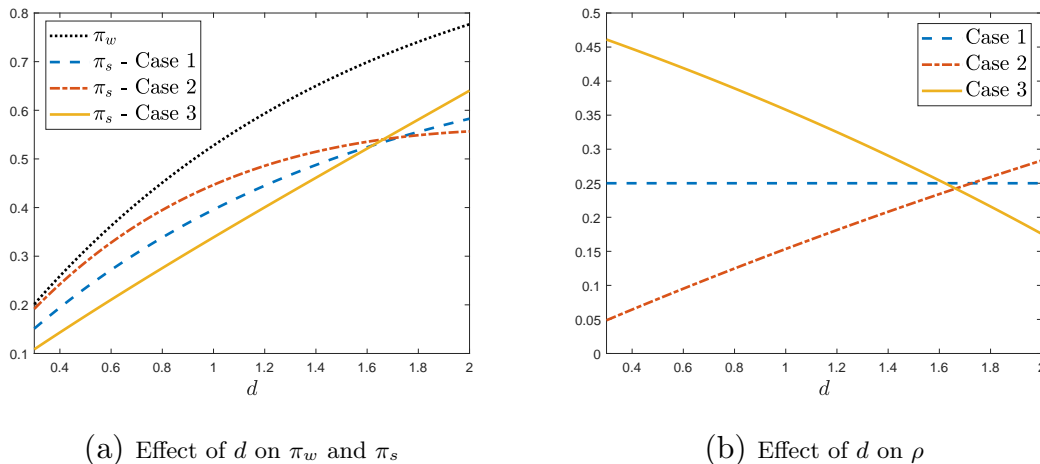
$$\rho(d) \triangleq \frac{\pi_w(d) - \pi_s(d)}{\pi_w(d)}. \quad (1)$$

We further assume that $1/\rho(d)$ is convex in d , which essentially requires data impact $\rho(\cdot)$ not to be too convex. This technical restriction is milder than log-concavity, and hence concavity.

Overall, when data impact $\rho(d)$ is constant in d , the effect of shirking on accuracy is independent of the data size on which the algorithm is trained. However, when the data impact increases (resp. decreases) in d , exerting effort increases the probability of high accuracy more (resp. less) with more data.

Figure 2 depicts different examples corresponding to these three regimes. In all examples, probability $\pi_w(\cdot)$ (the dotted black curve in Figure 2a) is the same, while probability $\pi_s(\cdot)$ takes different forms, inducing different properties for data impact ρ . In the first regime (Case 1 in Figure 2), the effect of shirking $\pi_w(d) - \pi_s(d)$ (when $\pi_s(d)$ is the blue dashed curve in Figure 2a) is directly proportional to the probability of high accuracy $\pi_w(d)$, which yields a constant data impact (the straight dashed line in Figure 2b). By contrast, the data impact increases in the second regime (the red dashed-dotted curve in Figure 2b) but decreases in the third one (the yellow plain curve in Figure 2b).

The monotonicity of the data impact is related to the monotone likelihood ratio property (MLRP), which is commonly assumed in the moral hazard literature. In our setup, the

Figure 2 The different regimes of the data impact.

Note. In all cases, the probability of high accuracy is $\pi_w(d) = 1 - \exp(-3d/4)$. We set $\pi_s(d) = \pi_w(d) \exp(-cd - y)$, and $\rho(d) = 1 - \exp(-cd - y)$, where the pair (c, y) is taken as $(0, \log 4/3)$ in Case 1 (constant ρ), $(1/6, 0)$ in Case 2 (increasing ρ) and $(-1/4, \log 2)$ in Case 3 (decreasing ρ).

MLRP property corresponds to $\pi_w(d)/\pi_s(d) \geq \bar{\pi}_w(d)/\bar{\pi}_s(d)$ for a given d and holds for all d since effort always improves accuracy $\pi_w(d) > \pi_s(d)$ (see Dai and Jerath 2019, for instance). Loosely speaking, the property guarantees that high accuracy is more indicative of high effort *ex-post*. The monotonicity of the data impact determines the magnitude of this effect *ex ante* for the amount of available data to train the algorithm.

2.3. Contracts

The firm faces a moral hazard problem in each period but do not have commitment power across periods. Thus, the firm needs to offer two different contracts, one per period. These contracts are nonetheless linked since the second contract determines the provider's expected payoff in the second period, which affects the provider's incentives to exert effort in the first period.

Accuracy realizations α_t , $t = 1, 2$ are contractable but effort is not. Thus, the contract in each period offers payments that depend on the public history and are contingent upon the accuracy realizations. Specifically, the history in Period 1a reduces to d_0 and the first contract consists of payments $x_{1\ell}(d_0)$ and $x_{1h}(d_0)$ that are made if $\alpha_1 = \alpha_\ell$ and $\alpha_1 = \alpha_h$, respectively. Given realization α_1 , the firm collects additional data, which yields size d_1 at the end of the period. The history in Period 2a is then (d_1, α_1) , and the second contract consists of payments $x_{2\ell}(d_1, \alpha_1)$ and $x_{2h}(d_1, \alpha_1)$, which again correspond to high and low

accuracy levels, respectively.² The provider is further protected by its limited liability; thus, $x_{1\ell}, x_{1h}, x_{2\ell}$ and x_{2h} are all non-negative.

2.4. The Firm's Problem

The firm's problem is to maximize the total expected profit, which is the expected revenue net of payments over both periods, subject to incentive compatibility constraints. We assume that the firm prefers the provider to exert effort in both periods and that the optimal price neither covers nor excludes the entire market. (These assumptions are made for the sake of simplicity; see, for instance, Laffont and Martimort 2009 and Feldman et al. 2018, Choudhary et al. 2005, respectively. We provide formal conditions for these assumptions in Appendix B).

We formulate this problem via backward induction starting from the second period (see Figure 1). Denote then by $J_{2b}(\alpha_2, d_1)$ the firm's optimal expected profit in Period 2b given accuracy α_2 and data size d_1 , such that

$$J_{2b}(\alpha_2, d_1) = \max_{\delta_2 \in [0, \nu]} R_\nu(\delta_2, \alpha_2) \quad (2)$$

The firm chooses the amount of collected data (or equivalently the price) so as to maximize the expected revenue in the current period. As there is no continuation, we refer to this problem as the *myopic* problem. In particular, size d_1 does not play any role in this problem, which corresponds to situations where both the AI Flywheel effect and the moral hazard problem are absent. We denote by δ^M the value of δ_2 that solves Problem (2).

Similarly, we denote by $J_{2a}(\alpha_1, d_1)$ the firm's optimal expected profit in Period 2a, given accuracy α_1 and data size d_1 , such that

$$J_{2a}(\alpha_\ell, d_1) = \max_{x_{2h}, x_{2\ell} \geq 0} \pi_w(d_1)[J_{2b}(\alpha_h, d_1) - x_{2h}] + \bar{\pi}_w(d_1)[J_{2b}(\alpha_\ell, d_1) - x_{2\ell}] \quad (3)$$

s.t.

$$\pi_w(d_1)x_{2h} + \bar{\pi}_w(d_1)x_{2\ell} - \kappa \geq \pi_s(d_1)x_{2h} + \bar{\pi}_s(d_1)x_{2\ell} \quad (4)$$

$$J_{2a}(\alpha_h, d_1) = J_{2b}(\alpha_h, d_1). \quad (5)$$

When $\alpha_1 = \alpha_\ell$, the firm needs to set payments such that the provider has enough incentives to exert effort, as formalized by incentive constraint (4). These payments are then deduced

² Recall that these last payments are only meaningful when $\alpha_1 = \alpha_\ell$, as no contract is required in the second period when $\alpha_1 = \alpha_h$.

from the firm's expected revenues in Period 2a. Here, data size d_1 affects the chance of improving accuracy in the next period via probabilities $\pi_e(\cdot)$, $e \in \{w, s\}$. When $\alpha_1 = \alpha_h$, recall that the firm does not need nor pay the provider. We thus refer to x_{2h}^* and $x_{2\ell}^*$ as the optimal payments solving Problem (3).

Moving to the first period, we denote by $J_{1b}(\alpha_1, d_0)$ the firm's optimal expected profit in Period 1b given accuracy α_1 and data size d_0 , such that

$$J_{1b}(\alpha_1, d_0) = \max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_1) + J_{2a}(\alpha_1, d_0 + \delta_1) \text{ for } \alpha_1 \in \{\alpha_h, \alpha_\ell\}. \quad (6)$$

In contrast to Problem (2), the firm needs to balance the revenues in the current period with the expected profit in the next one when the algorithm is of low accuracy. When $\alpha_1 = \alpha_\ell$, the choice of data δ_1 (or equivalently price) affects current revenues directly and future ones indirectly by increasing the dataset size to $d_0 + \delta_1$. We refer to δ^* as the optimal solution of Problem (6) for $\alpha_1 = \alpha_\ell$, and to p^* as the price that yields data size δ^* (see Section 2). When $\alpha_1 = \alpha_h$, no retraining is required and the optimal price is equal to the myopic price, as we make clear in Section 4.

We are now ready to define the overall firm's problem. Given initial data size d_0 , we denote by $J_{1a}(d_0)$ the optimal total expected profit in Period 1a, such that

$$J_{1a}(d_0) = \max_{x_{1h}, x_{1\ell} \geq 0} \pi_w(d_0)[J_{1b}(\alpha_h, d_0) - x_{1h}] + \bar{\pi}_w(d_0)[J_{1b}(\alpha_\ell, d_0) - x_{1\ell}] \quad (7)$$

s. t.

$$\pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[x_{1\ell} + J_p(d_0)] - \kappa \geq \pi_s(d_0)x_{1h} + \bar{\pi}_s(d_0)[x_{1\ell} + J_p(d_0)] \quad (8)$$

$$J_p(d_0) = \pi_w(d_0 + \delta^*)x_{2h}^* + \bar{\pi}_w(d_0 + \delta^*)x_{2\ell}^* - \kappa. \quad (9)$$

The firm faces a similar trade-off as in Period 2a. The difference, however, is in the incentive compatibility constraint (8). Indeed, the expected payments that the contract of the second period brings about affect the provider's incentives in the first period. Specifically, $J_p(\cdot)$ in (8) corresponds to the provider's expected continuation profit, which is equal to the expected optimal payments in the second period net of the effort cost; see (9). A key aspect of our setup is that the distribution of these future payments, $\pi_w(d_1)$, depends on the choice of δ_1 since $d_1 = d_0 + \delta_1$. Thus, the choice of δ_1 (or equivalently price) not only makes the tradeoff between present and future revenues as in (6) but also determines the intensity of the incentive issue.

3. First-Best Benchmark

We first consider the first-best setting in which the firm has the capability to develop the algorithm and does not face any incentive issues. The first-best problem then corresponds to Problem (7) without incentive constraints (8) and (4), but where the firm directly incurs cost κ . Again, we solve the first-best problem using backward induction.

Specifically, the firm's problem in Period 2b still corresponds to the myopic problem in (2). Straightforward calculations then the following result.

LEMMA 1. *Given accuracy α_2 and dataset size d_1 , the optimal collected data size δ^M and expected profit $J_{2b}(\alpha_2, d_1)$ are equal to $\delta^M = \nu \bar{F}(\phi^{-1}(0))$ and $J_{2b}(\alpha_2, d_1) = \alpha_2 \tau$, respectively, where $\tau \triangleq \phi^{-1}(0) \bar{F}(\phi^{-1}(0))$ and $\phi^{-1}(\cdot)$ is the inverse of virtual value function $\phi(\cdot)$.*

The proof of Lemma 1 are provided alongside all other proofs in Appendix A. The corresponding myopic price that yields data size δ^M is then equal to, per Section 2,

$$p^M = \alpha_2 \phi^{-1}(0),$$

and quantity τ is equal to $p^M(\delta^M/\nu)/\alpha_2$, which is the marginal revenue per unit of accuracy under optimal myopic pricing. Note that data size δ^M does not depend on accuracy, but optimal price p^M and profit J_{2b} do.

In period 2a, no contract is required at first-best, but the firm incurs the effort cost. The corresponding expected profit, J_{2a}^{FB} , is then given by equation (3) without payments but with cost κ . Given data size d_1 , expected profit J_{2a}^{FB} is thus equal to, where the future expected profit is equal to $\alpha_2 \tau$ per Lemma 1,

$$J_{2a}^{\text{FB}}(\alpha_\ell, d_1) = \pi_w(d_1) \alpha_h \tau + \bar{\pi}_w(d_1) \alpha_\ell \tau - \kappa \quad (10)$$

$$J_{2a}^{\text{FB}}(\alpha_h, d_1) = \alpha_h \tau. \quad (11)$$

The optimal expected profit J_{1b}^{FB} in Period 1b then corresponds to (6), where the future expected value is given by (10), i.e.,

$$J_{1b}^{\text{FB}}(\alpha_1, d_0) = \max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_1) + J_{2a}^{\text{FB}}(\alpha_1, d_0 + \delta_1) \text{ for } \alpha_1 \in \{\alpha_h, \alpha_\ell\}. \quad (12)$$

We denote by δ^{FB} the optimal value of δ_1 maximizing (12) for $\alpha_1 = \alpha_\ell$. We also refer to p^{FB} as the corresponding price that yields data size δ^{FB} . When $\alpha_1 = \alpha_h$, no training is required and the first-best price is equal to p^M , as we show later in this section.

Finally, no contract is required at first-best in Period 1a. The firm does not need to specify any payment but nonetheless incurs the effort cost (recall that we assume that exerting effort is always optimal; see Section 2.4). Given initial dataset size d_0 , the total expected profit at first-best is equal to

$$J_{1a}^{\text{FB}}(d_0) = \pi_w(d_0)J_{1b}^{\text{FB}}(\alpha_h, d_0) + \bar{\pi}_w(d_0)J_{1b}^{\text{FB}}(\alpha_\ell, d_0) - \kappa. \quad (13)$$

The following proposition characterizes the firm's optimal decision at first-best.

PROPOSITION 1. *The optimal solution to Problem (12) is unique and such that $\delta^{\text{FB}} \geq \delta^{\text{M}}$ if $\alpha_1 = \alpha_\ell$ and δ^{M} if $\alpha_1 = \alpha_h$.*

If the firm succeeds in developing a first algorithm of high accuracy ($\alpha_1 = \alpha_h$), no further improvement is necessary and the firm charges the optimal myopic price p^{M} , inducing δ^{M} over the remaining time horizon. If this accuracy is low ($\alpha_1 = \alpha_\ell$), however, the firm faces a tradeoff between maximizing revenues in the current period or acquiring additional data to leverage the AI Flywheel effect. In this case, the firm underprices with $p^{\text{FB}} < p^{\text{M}}$ and forfeits the optimal myopic revenue to collect more data, i.e., $\delta^{\text{FB}} > \delta^{\text{M}}$, increasing the probability of high accuracy and hence expected profit in the next period.

This scenario has further implications for the effect of the initial dataset on the firm's pricing strategy. Specifically, the firm may put more emphasis on increasing its initial dataset size before developing the algorithm, for instance, by purchasing or manually collecting additional data (see, e.g., Brown 2015 and Roh et al. 2019). The next proposition characterizes the ensuing effect on the firm's decisions.

PROPOSITION 2. *The first-best data size δ^{FB} is nonincreasing in initial data size d_0 .*

In other words, the first-best price p^{FB} increases with the initial data size d_0 . Per Proposition 1, the first-best benchmark δ^{FB} is always larger than myopic data size δ^{M} . Increasing d_0 allows for a higher $\pi_w(d_0)$ without affecting the first period revenue, thereby diminishing the need for deviating the price away from myopic price p^{M} .

Instead of acquiring a larger initial dataset, the firm may also consider designing a product that collects more data per user. This may be done, for instance, by using third-party services in the case of mobile and web applications or increasing the capacity of embedded sensors in the case of physical products (McGrath and Scanail 2013). In our setup, collecting more data per user corresponds to increasing volume rate ν , the effect of which is characterized by the next result.

PROPOSITION 3. *The first-best data size δ^{FB} is nondecreasing in data volume rate ν .*

In other words, as the volume rate increases, the firm collects even more data (δ^{FB} is nondecreasing). Recall that $\delta = \nu q$ and thus increasing ν provides an opportunity to increase data size δ to improve the algorithm in the next period while maintaining quantity q and hence revenues in the current period. Finally, note that parameters d_0 and ν have opposite effects on the firm's decisions.

4. Optimal Decisions of the Firm with Incentive Issues

We now characterize the optimal decisions of the firm when it needs to outsource the development of the algorithm. These decisions correspond to the amount of collected data δ^* (and the corresponding price p^*) as well as all payments $x_{1\ell}^*$, x_{1h}^* , $x_{2\ell}^*$, and x_{2h}^* . As in the previous sections, we study this problem using backward induction.

The optimization problem of J_{2a} in (3) corresponds to minimizing the expected payment while motivating the provider to exert effort. The following proposition (based on Laffont and Martimort 2009, Proposition 4.2) provides the corresponding optimal payments.

PROPOSITION 4. *Given dataset size d_1 , the unique optimal payments for Problem (3) are $x_{2\ell}^* = 0$ and $x_{2h}^* = \kappa / [\pi_w(d_1) - \pi_s(d_1)]$. The firm's optimal expected profit is then*

$$J_{2a}(\alpha_\ell, d_1) = \pi_w(d_1)\tau\alpha_h + \bar{\pi}_w(d_1)\tau\alpha_\ell - \frac{\kappa}{\rho(d_1)} \text{ and } J_{2a}(\alpha_h, d_1) = \tau\alpha_h. \quad (14)$$

where $\kappa/\rho(d_1)$ is the expected payment to the provider.

MLRP ensures that under the optimal contract, realizations of higher value (α_h in our setup) are more rewarded ($x_{2h}^* > x_{2\ell}^*$). Thus, MLRP concerns the ex post payments to the provider. In contrast, Proposition 4 shows that data impact $\rho(\cdot)$ determines the ex ante provider's rent $\kappa/\rho(d_1)$.

The optimal expected continuation profit $J_{1b}(\alpha_1, d_0)$ in Problem (6) is then obtained by using expected profit $J_{2a}(\alpha_1, d_0 + \delta)$ from (14) in Proposition 4. We show next that the optimal data size solving Problem (6) and hence the corresponding optimal price are unique.

PROPOSITION 5. *The optimal solution to Problem (6) is unique and equal to δ^{M} if $\alpha_1 = \alpha_h$.*

If the first version of the algorithm is already highly accurate ($\alpha_1 = \alpha_h$), no further improvement is necessary and the firm does not need to deviate from the first-best price, which is also the myopic price ($p^{\text{FB}} = p^{\text{M}}$ from Proposition 5). In the next section, we explore in detail the deviations from first-best that a low accuracy ($\alpha_1 = \alpha_\ell$) creates.

Finally, the next result characterizes the optimal payments in Period 1a.

PROPOSITION 6. *Given initial data size d_0 , the unique optimal payments for Problem (7) are*

$$x_{1\ell}^* = 0 \text{ and } x_{1h}^* = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + \left(\frac{\kappa}{\rho(d_0 + \delta^*)} - \kappa \right).$$

Contrary to the payments in Period 2, the firm now needs to account for the provider's future expected profits. Specifically, optimal payment x_{1h}^* corresponds to i) bonus payment taken at d_0 but augmented by ii) the provider's rent of the second period net of effort cost because in our setup, the firm cannot easily replace the provider across periods. In this sense, the second term of x_{1h}^* captures the cost due to the scarcity of AI service providers in the market.

5. Overpricing, Underpricing and Optimal Data Collection

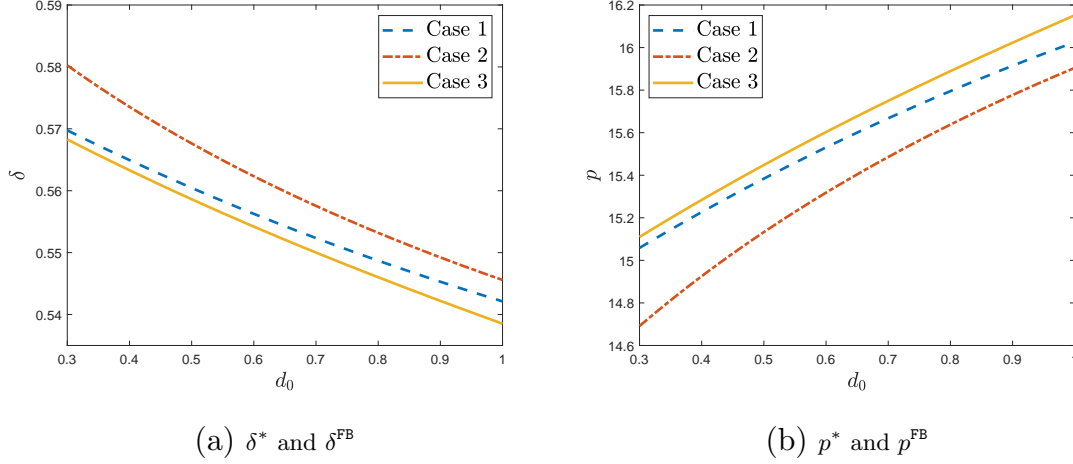
To alleviate the moral hazard issue, the firm may need to incur costly deviations away from first-best decisions. The next result shows that the monotonicity of data impact $\rho(\cdot)$ is sufficient to determine when the firm overprices and when it underprices.

THEOREM 1. *We have*

1. *if $\rho(\cdot)$ is constant, then $\delta^* = \delta^{\text{FB}}$,*
2. *if $\rho(\cdot)$ is increasing, then $\delta^* \geq \delta^{\text{FB}}$,*
3. *if $\rho(\cdot)$ is decreasing, then $\delta^* \leq \delta^{\text{FB}}$.*

In essence, the monotonicity of the data impact defines three distinct regimes. If the relative impact of shirking is independent of the dataset size (ρ is constant), then no deviation from first-best is required. By contrast, if this impact intensifies with more data (ρ is increasing), the firm needs to underprice $p^* \leq p^{\text{FB}}$ in order to collect more data $\delta^* \geq \delta^{\text{FB}}$. Finally, the firm needs to overprice $p^* \geq p^{\text{FB}}$ when the relative impact of shirking diminishes with more data (ρ is decreasing).

Figure 3 depicts examples of these three regimes for different initial dataset sizes (d_0). The data impacts corresponding to these examples are depicted in Figure 2b. The resulting data sizes δ^* and optimal prices p^* are depicted in Figures 3a and 3b, respectively.

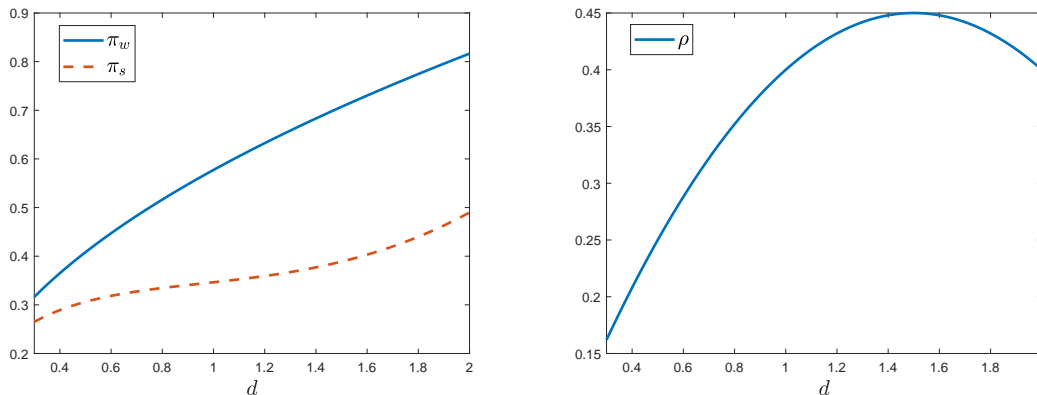
Figure 3 Optimal data and pricing decisions of the firm.

Note. The data impacts for Case 1, 2, and 3 are given by the corresponding cases in Figure 2b. For Case 1, $\delta^{\text{FB}} = \delta^*$ and $p^{\text{FB}} = p^*$ are in Figure 3a and Figure 3b, respectively.

In the first regime, the data impact is constant and the net effect of shirking $\pi_w(d) - \pi_s(d)$ is proportional to the probability of high accuracy $\pi_w(d)$ (see Figure 2, Case 1). In this case, the dataset size does not affect the intensity of the incentive issues. The firm only faces the tradeoffs that the AI Flywheel effect brings about and thus charges the first-best price (but still incurs the payments to the provider). This corresponds to the blue dashed curve in Figure 3, which depicts the firm's decisions in both the first regime and first-best.

In the second regime, the data impact is increasing and the net effect of shirking increases faster than the probability of high accuracy (see Figure 2, Case 2). This regime is depicted by the red dashed-dotted curve in Figure 3. In this case, high accuracy is more indicative of efforts at higher data volumes, and the rent is decreasing in data size δ . Thus, the firm underprices with $p^* < p^{\text{FB}}$ (the red dashed-dotted curve is below the blue dashed one in Figure 3b) in order to collect more data with $\delta^* > \delta^{\text{FB}}$ (the red dashed-dotted curve is above the blue dashed one in Figure 3a). In this sense, the incentive issues induce the firm to leverage the AI Flywheel effect even more. In particular, the expected revenue in the first period is lower than first-best due to incentive issues, but the expected revenues in the second period are higher since with more data, the probability of high accuracy is higher.

Finally, the data impact is decreasing (see Figure 2, Case 3) in the third regime, which is depicted by the yellow plain curve in Figure 3. In this case, high accuracy is less indicative of efforts at higher data levels and the rent is increasing in data size δ . Thus, the firm overprices with $p^* > p^{\text{FB}}$ so as to collect less data with $\delta^* < \delta^{\text{FB}}$. The firm's expected revenues

Figure 4 Unimodal data impact.(a) Effect of d on π_w and π_s (b) Effect of d on ρ

Note. Probabilities $\pi_w(d) = \sqrt{d/3}$ and $\pi_s(d) = \pi_w(d)(1 - 0.2(3 - d)d)$ are depicted in Figure 4a and induce data impact $\rho(d) = 0.2(3 - d)d$ is depicted in Figure 4b.

in the last period decrease compared to first-best since the probability of high accuracy is lower. The revenues in the first period, however, might actually increase compare to first-best. This situation happens, for instance, when $p^M > p^* > p^{FB}$. Overall, the incentive issues prevent the firm from fully leveraging the AI Flywheel effect in this regime.

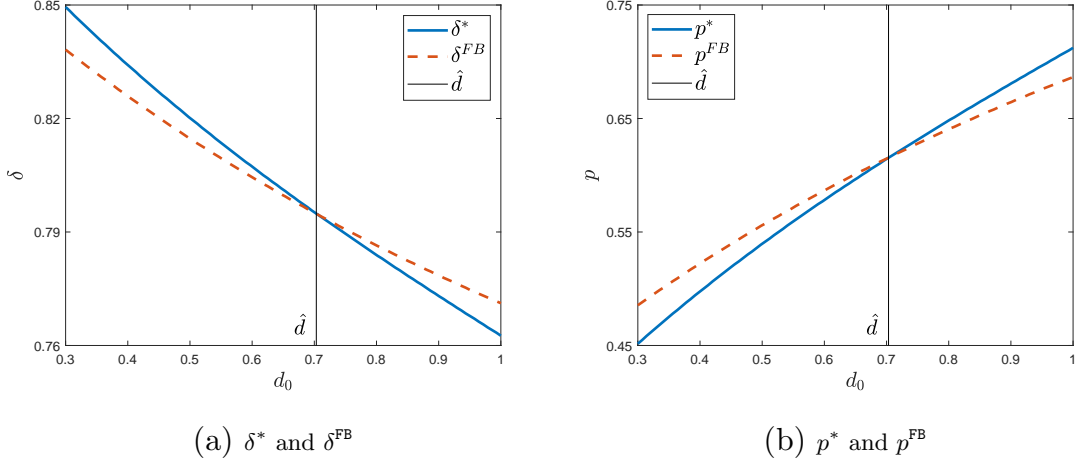
An important consequence of Theorem 1 is that the initial dataset does not determine whether the firm overprices or underprices, as long as the data impact is monotone. This is not the case, however, when the data impact is unimodal, as demonstrated by the following corollary.

COROLLARY 1. *If data impact $\rho(\cdot)$ is unimodal, then threshold \hat{d} exists such that*

1. *if $d_0 = \hat{d}$, then $\delta^* = \delta^{FB}$,*
2. *if $d_0 \leq \hat{d}$, then $\delta^* \geq \delta^{FB}$,*
3. *if $d_0 \geq \hat{d}$, then $\delta^* \leq \delta^{FB}$.*

Thus, the size of the initial dataset determines whether the firm underprices or overprices when the data impact is unimodal. Specifically, the firm underprices when the dataset is small but overprices otherwise. In this sense, the size of the initial dataset may reverse the firm's overall pricing strategy.

The data impact is unimodal when the relative effect of shirking is highest for datasets of medium size. Figure 4 provides an example of probabilities π_w and π_s (Figure 4a) that yield a unimodal data impact (Figure 4b), which corresponds to the situation where the

Figure 5 Effect of the initial dataset on the optimal data collection and pricing decisions when ρ is unimodal.

Note. The corresponding data impact is depicted in Figure 4b.

sheer volume of data, rather than the algorithm's refinements, determine the algorithm's accuracy for large datasets, while the lack of data limits the effect of any effort levels for small datasets.

Figures 5a and 5b depict the effect of initial size d_0 on the firm's data collection and pricing strategies, respectively. The data impact corresponding to this example is depicted in Figure 4b. The optimal data size δ^* (resp. price p^*) is above (resp. below) the collected data size δ^{FB} at first-best when d_0 is less than threshold $\hat{d} \approx 0.7$ and below it otherwise.

6. Impact of the Firm's Initial Dataset

The previous results show that the data impact's monotonicity determines the firm's pricing and data collection strategies. These results also suggest that the size of the initial dataset sometimes plays a key role in the firm's decisions. We explore this role further below.

In our setup, initial dataset size d_0 determines the firm's starting point in the virtuous cycle of the AI Flywheel effect. The firm can increase this dataset by manually collecting more data or purchasing existing datasets. This scenario, however, affects the firm's decisions, as shown in the following result.

PROPOSITION 7. *The optimal data size δ^* is nonincreasing in d_0 .*

Hence, the larger the initial dataset is, the higher the price charged by the firm and thus the less data are collected across periods. Interestingly, this result does not depend on the monotonicity of the data impact, as illustrated by Figures 3 and 5.

Further, recall that an increase of initial size d_0 also increases collected data δ^* (see Proposition 2). Thus, whether an increase in the initial dataset exacerbates or reduces the distortions away from first-best remains unknown in general. Nonetheless, the next result provides different sufficient conditions, under which absolute distortion $|\delta^* - \delta^{\text{FB}}|$ may increase or decrease.

PROPOSITION 8. *The absolute deviation away from first-best is such that*

- $|\delta^* - \delta^{\text{FB}}|$ is nonincreasing in d if $\rho(\cdot)$ is increasing and $\pi_w''(\cdot)$ is nonincreasing.
- $|\delta^* - \delta^{\text{FB}}|$ is nondecreasing in d if $\rho(\cdot)$ is decreasing and $\pi_w''(\cdot)$ is nondecreasing.

Hence, under third-order conditions on probability π_w , an increase in the initial data size reduces the distortions away from first-best (p^* and δ^* approaches p^{FB} and δ^{FB} , respectively) when the data impact is increasing but exacerbates them when the data impact is decreasing (p^* and δ^* move away from p^{FB} and δ^{FB} , respectively). The monotonicity of the data impact determines whether or not δ^* is larger than δ^{FB} (see Theorem 1). Furthermore, Propositions 2 and 7 show that both δ^{FB} and δ^* are decreasing in initial size d_0 . The condition on probability of high accuracy π_w then ensures that optimal data size δ^* decreases faster than δ^{FB} in d_0 . When the data impact is decreasing (resp. increasing), this condition states that the probability becomes more (resp. less) concave with data.

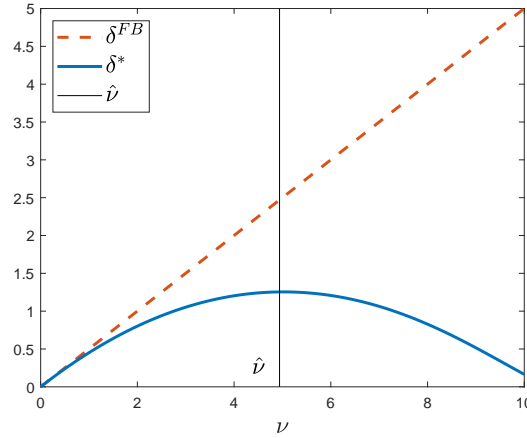
7. Impact of the Data Volume Rate

Besides using its initial dataset, the firm may also seek to design a product that collects more usage data per user. In our setup, this scenario corresponds to improving volume rate ν . Without incentive issues, a higher volume rate induces the firm to collect even more data under the AI Flywheel effect per Proposition 3. The need to mitigate incentive issues, however, sometimes reverses this effect and pushes the firm to collect *less* data when the volume rate is higher. The next theorem, one of our main results, formalizes this finding.

THEOREM 2. *We have the following:*

1. *If $\rho(\cdot)$ is increasing, δ^* is nondecreasing in ν .*
2. *If $\rho(\cdot)$ is decreasing, a unique threshold $\hat{\nu}$ exists such that δ^* is nondecreasing in ν if $\nu \leq \hat{\nu}$ and is nonincreasing otherwise.*

When the volume rate increases, the firm has an incentive to collect more data to benefit more from the AI Flywheel effect, as discussed in the first-best benchmark (see Proposition

Figure 6 Effect of volume rate on the optimal data collection decision.

Note. Probabilities generating this example are $\pi_w(d) = 1 - \exp(-2d)$, $\pi_s(d) = \pi_w(d)(1 - 1/(d+1))$ and the data impact is $\rho(d) = 1/(1+d)$.

3). When the data impact is increasing, collecting more data also reduces the intensity of the moral hazard problem and hence the agency costs in the next period. Both effects are aligned in this case, and the firm increases δ^* as a result. By contrast, collecting more data intensifies the moral hazard problem in the next period when the data impact is decreasing. The firm then faces a tradeoff between leveraging the AI Flywheel effect or reducing future agency costs. When the volume rate is small, the amount of collected data $\delta = \nu q$ remains small for any quantity q , which yields low agency costs. As the volume rate increases, boosting the AI Flywheel effect dominates the increases in agency costs, and data size δ^* increases as a result. When the volume rate becomes large enough (i.e., when $\nu > \hat{\nu}$), however, the agency costs dominates the revenues due to the AI Flywheel effect. The firm then focuses on reducing these agency costs by decreasing δ^* .

Figure 6 illustrates the second item of Theorem 2 with an example. While the first-best benchmark δ^{FB} (the red dashed curve in Figure 6) increases as suggested in Proposition 3, the optimal data δ^* demonstrates a nonmonotone behavior in the volume rate ν . Threshold $\hat{\nu}$ corresponds to the straight black line in Figure 6), with $\hat{\nu} \cong 4.9$.

8. Concluding Remarks

This paper proposes a simple dynamic framework to study how firms that outsource the development of their ML algorithm can leverage the AI Flywheel effect. Our setup accounts for the three main features of this problem: i) the tradeoff between improving algorithms' accuracy and maximizing revenues due to the AI Flywheel effect, ii) the need to manage

the incentive issues that outsourcing the algorithm brings about, and iii) the interaction between the amount of data on which the algorithm is trained and the efficacy of the provider's effort. We further introduce the notion of data impact as a framework to represent the interaction between data and effort.

Taken together, our results identify three different regimes, which depend on the nature of the data impact. These regimes determine whether the firm overprices or underprices and the impact of both the starting dataset and the volume rate on the firm's decisions. In particular, when the data impact decreases, the firm sometimes acquires less data overall if the firm increases the data volume that its product collects. This effect stems directly from the need to manage the incentive issues that outsourcing the algorithm creates.

These results further provide predictions that future work can empirically test. In particular, given that the existing literature points to the importance of collecting a large amount of data over improving algorithms (Banko and Brill 2001, Halevy et al. 2009), we expect the data impact to decrease in many practical contexts. Our work thus provides theoretical support for the hypothesis that firms set higher prices for a new AI product when the product's algorithm is outsourced (per Theorem 1). In addition, we predict that a significant increase in the product's capacity to generate usage data induces firms to collect less data to improve the algorithm (per Theorem 2).

More generally, we believe that our work opens up new research directions and questions for the management of data-driven business models. Specifically, our paper considers a problem in which the provision of data interacts with incentive issues. Indeed, the key aspect of our setup is that the principal can regulate the intensity of the moral hazard problem she faces by controlling (through pricing in our setting) the data to which the agent has access. We believe that this interaction between data and incentives is present in many other contexts than the outsourcing of the AI Flywheel effect. Our framework offers a fruitful starting point to model and study these issues.

To the best of our knowledge, these results provide the first insights on how firms can leverage the AI Flywheel effect. In addition, ours is the first paper to consider the problem of contracting ML algorithms. Given the shortage of data scientists and the growing outsourcing industry in this domain, we expect the issue to gain importance in the coming years.

References

- Acemoglu D, Makhdoui A, Malekian A, Ozdaglar A (2019) Too much data: Prices and inefficiencies in data markets. Technical report, National Bureau of Economic Research.
- Aflaki A, Feldman P, Swinney R (2019) Becoming strategic: Endogenous consumer time preferences and multiperiod pricing. *Forthcoming, Operations Research* .
- Anderson Jr EG, Parker GG (2013) Integration and cospecialization of emerging complementary technologies by startups. *Production and Operations Management* 22(6):1356–1373.
- Banko M, Brill E (2001) Scaling to very very large corpora for natural language disambiguation. *Proceedings of the 39th annual meeting on association for computational linguistics*, 26–33 (Association for Computational Linguistics).
- Bimpikis K, Crapis D, Tahbaz-Salehi A (2019) Information sale and competition. *Management Science* 65(6):2646–2664.
- Brown MS (2015) When and where to buy consumer data (and 12 companies who sell it). <https://www.forbes.com/sites/metabrown/2015/09/30/when-and-where-to-buy-consumer-data-and-12-companies-who-sell-it/#413d4b133285> [Online; accessed 10-March-2020].
- Bursztein E (2018) Attacks against machine learning ? an overview. elie.net/blog/ai/attacks-against-machine-learning-an-overview/ [Online; accessed 27-January-2020].
- Cengiz K (2015) 10 startups that have become a success with outsourced development. <https://perfectial.com/blog/?s=10+Startups+That+Have+Become+A+Success+With+Outsourced+Development&search-type=normal> [Online; accessed 17-January-2020].
- Choudhary V, Ghose A, Mukhopadhyay T, Rajan U (2005) Personalized pricing and quality differentiation. *Management Science* 51(7):1120–1130.
- Dai T, Jerath K (2019) Salesforce contracting under uncertain demand and supply: Double moral hazard and optimality of smooth contracts. *Marketing Science, Forthcoming* .
- de Véricourt F, Gromb D (2018) Financing capacity investment under demand uncertainty: An optimal contracting approach. *Manufacturing & Service Operations Management* 20(1):85–96.
- de Véricourt F, Gromb D (2019) Financing capacity with stealing and shirking. *Management Science* 65(11):5128–5141.
- Deshpande I (2019) What is customer data? definition, types, collection, validation and analysis. <https://www.martechadvisor.com/articles/data-management/customer-data-definition-types-collection-validation-analysis-martech101/> [Online; accessed 10-March-2020].
- Feldman P, Papanastasiou Y, Segev E (2018) Social learning and the design of new experience goods. *Management Science* 65(4):1502–1519.

- Golden K (2017) Deere to advance machine learning capabilities in acquisition of blue river technology. www.deere.com/en/our-company/news-and-announcements/news-releases/2017/corporate/2017sep06-blue-river-technology/ [Online; accessed 27-January-2020].
- Halevy A, Norvig P, Pereira F (2009) The Unreasonable Effectiveness of Data. *IEEE Intelligent Systems* 24(2):8–12, ISSN 1541-1672, URL <http://dx.doi.org/10.1109/MIS.2009.36>.
- Jo J, Bengio Y (2017) Measuring the tendency of cnns to learn surface statistical regularities. *arXiv preprint arXiv:1711.11561* .
- Kendra A, Kumar S, O’Brien D, Shankar R, Snover J, Viljoen S (2019) Failure modes in machine learning. <https://docs.microsoft.com/en-us/security/failure-modes-in-machine-learning> [Online; accessed 17-January-2020].
- Kräkel M, Schöttner A (2016) Optimal sales force compensation. *Journal of Economic Behavior & Organization* 126:179–195.
- Laffont JJ, Martimort D (2009) *The theory of incentives: the principal-agent model* (Princeton university press).
- Lipton ZC (2016) The mythos of model interpretability. *arXiv preprint arXiv:1606.03490* .
- Loertscher S, Marx LM (2019) Digital monopolies: Privacy protection or price regulation? .
- McGrath MJ, Scanaill CN (2013) Key sensor technology components: hardware and software overview. *Sensor Technologies*, 51–77 (Springer).
- Mehta S, Dawande M, Janakiraman G, Mookerjee V (2019) How to sell a dataset? pricing policies for data monetization. *Pricing Policies for Data Monetization (February 12, 2019)* .
- Miller R (2016) At amazon the flywheel effect drives innovation. techcrunch.com/2016/09/10/at-amazon-the-flywheel-effect-drives-innovation/ [Online; accessed 27-January-2020].
- Ng A (2018) Ai transformation playbook. www.coursera.org/learn/ai-for-everyone/lecture/kSWz6/ai-transformation-playbook-part-2 [Online; accessed 27-January-2020].
- Nicolaus Henke J, Bughin, Chui M, Manyika J, Saleh T, Wiseman B, Sethupathy G (2016) The age of analytics: competing in a data-driven world.
- Research Nester Pvt Ltd (2019) Data analytics outsourcing market.
- Ribeiro MT, Singh S, Guestrin C (2016) Why should i trust you?: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (ACM).
- Roh Y, Heo G, Whang SE (2019) A survey on data collection for machine learning: a big data-ai integration perspective. *IEEE Transactions on Knowledge and Data Engineering* .
- Sarikaya R (2019) How alexa learns. <https://blogs.scientificamerican.com/observations/how-alexa-learns/> [Online; accessed 27-January-2020].

- Schmitz PW (2005) Allocating control in agency problems with limited liability and sequential hidden actions. *RAND Journal of Economics* 318–336.
- Schöttner A (2016) Optimal sales force compensation in dynamic settings: Commissions vs. bonuses. *Management Science* 63(5):1529–1544.
- Swinney R, Cachon GP, Netessine S (2011) Capacity investment timing by start-ups and established firms in new markets. *Management Science* 57(4):763–777.
- Tanrisever F, Erzurumlu SS, Joglekar N (2012) Production, process investment, and the survival of debt-financed startup firms. *Production and Operations Management* 21(4):637–652.
- Trautman E (2018) The virtuous cycle of ai products. <https://www.eriktrautman.com/posts/the-virtuous-cycle-of-ai-products> [Online; accessed 27-January-2020].
- Yu M, Debo L, Kapuscinski R (2015) Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *Management Science* 62(2):410–435.

Appendix

A. Proof of Results

Proof of Lemma 1. Recall the definition $R_\nu(\delta, \alpha) = \alpha\delta/\nu F^{-1}((1 - \delta/\nu))$ in Section 2, and F is the c.d.f. of the standard uniform distribution. Therefore, it can be verified that $R_\nu(\delta, \alpha)$ is concave in δ , and the following first-order condition is sufficient for optimality.

$$\frac{1}{\nu} F^{-1}\left(1 - \frac{\delta}{\nu}\right) - \frac{\delta}{\nu^2} \frac{1}{f\left(F^{-1}\left(1 - \frac{\delta}{\nu}\right)\right)} = 0. \quad (15)$$

Let $\xi = F^{-1}\left(1 - \frac{\delta}{\nu}\right)$, hence $\delta/\nu = \bar{F}(\xi)$. We first multiply both sides of the equality with ν and then use this new notation ξ . Using the definition of the virtual value function ϕ , we have

$$\phi(\xi) = 0. \quad (16)$$

Because the virtual value function of the uniform distribution is increasing, and crosses 0 at $1/2$, we conclude that the unique optimal solution δ^M to Problem (2) is equal to $\nu\bar{F}(\phi^{-1}(0))$. Evaluating the objective function at the optimal solution and using the fact that $\tau = \phi^{-1}(0)\bar{F}(\phi^{-1}(0))$, we conclude that

$$J_{2b}(\alpha_2, d_1) = \alpha_2\tau.$$

Proof of Proposition 1. Using the expected profit J_{2a}^{FB} in (10), we write Problem (12) as follows.

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_1) + 1\{\alpha_1 = \alpha_\ell\} [\pi_w(d_0 + \delta_1)\alpha_h\tau + \bar{\pi}_w(d_0 + \delta_1)\alpha_\ell\tau - \kappa] \quad (17)$$

where $1\{\cdot\}$ is the indicator function.

If $\alpha_1 = \alpha_h$, the objective function in (17) becomes $R_\nu(\delta_1, \alpha_h)$. Lemma 1 implies the unique optimal solution in this case is δ^M . Otherwise, the optimal solution δ^{FB} solves the following problem

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell)\tau + \alpha_\ell\tau - \kappa \quad (18)$$

Because π_w , R_ν are concave and $d_0 + \delta_1$ is a linear function of δ_1 , the objective function in (18) is concave in δ_1 . This implies that the optimal solution δ^{FB} is unique.

We next prove $\delta^{\text{FB}} > \delta^{\text{M}}$. Recall that δ^{M} is the unique optimal solution to $\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell)$ and $R_\nu(\delta_1, \alpha_\ell)$ is concave (see Lemma 1), and $\pi_w(d_0 + \delta_1)$ is increasing in δ_1 . Marginal revenue of $R_\nu(\delta_1, \alpha_\ell)$ at $\delta_1 = \delta^{\text{M}}$ is 0 while $\pi'_w(d_0 + \delta^{\text{M}})$ is positive. Therefore, it follows that

$$\left[\frac{\partial R_\nu(\delta_1, \alpha_\ell)}{\partial \delta_1} + \frac{\partial \pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell)\tau}{\partial \delta_1} \right] \Big|_{\delta_1 = \delta^{\text{M}}} > 0$$

Because the objective function in (18) is concave in δ_1 , its derivative with respect to δ_1 is decreasing and equal to 0 at $\delta_1 = \delta^{\text{FB}}$. (Note that we consider interior solutions for δ^{FB} , see Proposition 10 in Appendix B.)

These observations imply that $\delta^{\text{FB}} > \delta^{\text{M}}$. Q.E.D.

Proof of Proposition 2. Proposition 1 implies that δ^{FB} satisfies the following first-order condition.

$$\left[\frac{\partial R_\nu(\delta_1, \alpha_\ell)}{\partial \delta_1} + \frac{\partial \pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell)\tau}{\partial \delta_1} \right] \Big|_{\delta_1 = \delta^{\text{FB}}} = 0 \quad (19)$$

Note here that the left-hand side of (19) is a continuously differentiable function of δ^{FB} because $R_\nu(\delta, \alpha_\ell)$ and $\pi_w(d_0 + \delta)$ are twice-differentiable with a continuous second-order derivative w.r.t. δ , thus we can use the implicit function theorem. Using the implicit function theorem, we know $\frac{\partial \delta^{\text{FB}}}{\partial d_0}$ has the same sign with $\frac{\partial^2 \pi_w(d_0 + \delta_1)}{\partial d_0 \partial \delta_1} \Big|_{\delta_1 = \delta^{\text{FB}}}$. Because π_w is concave and $d_0 + \delta^{\text{FB}}$ is linear in d_0 , it follows that $\frac{\partial \delta^{\text{FB}}}{\partial d_0} \leq 0$. Q.E.D.

Proof of Proposition 3. As in the proof of Proposition 2, we use the same first-order condition and the implicit function theorem to prove this result. Differently, $\frac{\partial \delta^{\text{FB}}}{\partial \nu}$ has the same sign with

$$\frac{\partial^2 R_\nu(\delta_1, \alpha_\ell)}{\partial \nu \partial \delta_1} \Big|_{\delta_1 = \delta^{\text{FB}}} = \frac{\alpha_\ell}{\nu^2} \left(-1 + \frac{4\delta^{\text{FB}}}{\nu} \right)$$

because we consider monotonicity with respect to ν . Recall that we know $\delta^{\text{FB}} \geq \delta^{\text{M}}$ from Proposition 1 and $\delta^{\text{M}} = \nu \bar{F}(\phi^{-1}(0)) = \nu/2$ from Lemma 1. Therefore, δ^{FB} is nondecreasing in ν .

Proof of Proposition 4. Following Proposition 4.2 in Laffont and Martimort (2009, p. 157), we obtain the optimal payments for Problem (3) $x_{2\ell}^* = 0$ and $x_{2h}^* = \kappa / [\pi_w(d_1) - \pi_s(d_1)]$. Following Lemma 1, we have $J_{2b}(\alpha_2, d_1) = \alpha_2 \tau$. Combining these, we get

$$J_{2a}(\alpha_\ell, d_1) = \pi_w(d_1)\tau\alpha_h + \bar{\pi}_w(d_1)\tau\alpha_\ell - \frac{\kappa}{\rho(d_1)}.$$

If $\alpha_1 = \alpha_h$, there are no payments so $J_{2a}(\alpha_h, d_1) = J_{2b}(\alpha_h, d_1) = \tau\alpha_h$. Q.E.D.

Proof of Proposition 5. Using the expected profit J_{2a} derived in (14), we write Problem (6) as follows.

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_1) + 1\{\alpha_1 = \alpha_\ell\} \left[\pi_w(d_0 + \delta_1)\alpha_h\tau + \bar{\pi}_w(d_0 + \delta_1)\alpha_\ell\tau - \frac{\kappa}{\rho(d_0 + \delta_1)} \right] \quad (20)$$

If $\alpha_1 = \alpha_h$, it is straightforward to see from (20) that the unique optimal solution to Problem (6) is δ^{M} .

Otherwise, the optimal solution δ^* solves the following problem

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell)\tau - \frac{\kappa}{\rho(d_0 + \delta_1)} + \alpha_\ell\tau \quad (21)$$

Because π_w , R_ν and $-1/\rho$ are concave, and $d_0 + \delta_1$ is a linear function of δ_1 , the objective function in (21) is concave in δ_1 . This implies that the optimal solution δ^* is unique. Note also that we consider interior solutions for δ^* , see Proposition 10 in Appendix B. Q.E.D.

Proof of Proposition 6. In Period 1a, the optimization problem of the firm is

$$J_{1a}(d_0) = \pi_w(d_0)J_{1b}(\alpha_h, d_0) + \bar{\pi}_w(d_0)J_{1b}(\alpha_\ell, d_0) - \min_{x_{1h}, x_{1\ell} \geq 0} \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)x_{1\ell}$$

st. (8), (9)

Using the optimal payments x_{2h}^* and $x_{2\ell}^*$ derived in Proposition 4, we evaluate the expected continuation profit $J_p(d_0)$ of the provider, and $J_p(d_0) = \frac{\kappa}{\rho(d_0 + \delta^*)} - \kappa$. Characterizing $J_p(d_0)$, we reduce this problem to a standard principal-agent model where the cost of effort is $\kappa + [\pi_w(d_0) - \pi_s(d_0)]J_p(d_0)$. Therefore, Proposition 4.2 in Laffont and Martimort (2009, p. 157) implies that the optimal solution of the optimization problem in $J_{1b}(d_0)$ is $x_{1\ell}^* = 0$ and $x_{1h}^* = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + \frac{\kappa}{\rho(d_0 + \delta^*)} - \kappa$. Q.E.D.

Proof of Theorem 1. Note that if $\rho(d)$ is constant, then the objective function in Problem (12) and the one in Problem (6) for $\alpha_1 = \alpha_\ell$ are different from each other by a constant. Therefore, their optimal solutions are the same, i.e., $\delta^{\text{FB}} = \delta^*$.

Next, assume that $\rho(d)$ is decreasing in d . We prove this item by contradiction. Assume that $\delta^* < \delta^{\text{FB}}$. Fix d_0 and ν . The following condition is satisfied by δ^{FB} because δ^{FB} is the unique optimal solution of $\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \tau\pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell)$ (see Proposition 1).

$$R_\nu(\delta^{\text{FB}}, \alpha_\ell) + \tau\pi_w(d_0 + \delta^{\text{FB}})(\alpha_h - \alpha_\ell) \geq R_\nu(\delta, \alpha_\ell) + \tau\pi_w(d_0 + \delta)(\alpha_h - \alpha_\ell), \forall \delta \in [0, \nu] \quad (22)$$

Because the data impact $\rho(\cdot)$ is decreasing, the term $-\frac{\kappa}{\rho(d_0 + \delta)}$ is an increasing function of δ . Using the assumption of contradiction, we obtain

$$-\frac{\kappa}{\rho(d_0 + \delta^{\text{FB}})} > -\frac{\kappa}{\rho(d_0 + \delta^*)} \quad (23)$$

Inequalities (22) and (23) imply that the objective function $R_\nu(\delta_1, \alpha_\ell) + \tau\pi_w(d_0 + \delta_1)(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d_0 + \delta_1)}$ evaluated at δ^{FB} is strictly larger than the value obtained by evaluating the same at δ^* . Therefore, the condition $\delta^* < \delta^{\text{FB}}$ contradicts with the fact that δ^* is the optimal solution (see Proposition 5). The last item of the theorem can be proved by following the same steps with $\delta^* > \delta^{\text{FB}}$. Q.E.D.

Proof of Corollary 1. We prove this result in two steps. Fix α_1 and ν . We first prove the following listed items. In the second step, we prove that $m(d_0) \triangleq d_0 + \delta^{\text{FB}}$ is an increasing function of d_0 .

Let \tilde{d} be the peak point of $\rho(\cdot)$.

1. if d_0 is such that $d_0 + \delta^{\text{FB}} = \tilde{d}$, then $\delta^* = \delta^{\text{FB}}$.
2. if d_0 is such that $d_0 + \delta^{\text{FB}} \leq \tilde{d}$, then $\delta^* \geq \delta^{\text{FB}}$.
3. if d_0 is such that $d_0 + \delta^{\text{FB}} \geq \tilde{d}$, then $\delta^* \leq \delta^{\text{FB}}$.

Step 1. In this step, we follow a procedure similar to the proof of Theorem 1 for the second and third items: proof by contradiction. Assume that $d_0 + \delta^{\text{FB}} \leq \tilde{d}$ and $\delta^* < \delta^{\text{FB}}$. Then the following inequality holds.

$$R_\nu(\delta^{\text{FB}}, \alpha_\ell) + \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta^{\text{FB}}) - \frac{\kappa}{\rho(d_0 + \delta^{\text{FB}})} \geq R_\nu(\delta^*, \alpha_\ell) + \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta^*) - \frac{\kappa}{\rho(d_0 + \delta^*)}$$

Here the left-hand side is larger because i) the sum of first two terms at the left-hand side is maximized at δ^{FB} , and ii) $-\frac{\kappa}{\rho(d_0 + \delta^{\text{FB}})} > -\frac{\kappa}{\rho(d_0 + \delta^*)}$ when $d_0 + \delta^{\text{FB}} \leq \tilde{d}$ and $\delta^* < \delta^{\text{FB}}$ for unimodal ρ with peak point \tilde{d} . This inequality contradicts with the fact that δ^* is the optimal solution as shown in Proposition 5. Therefore

it follows that $\delta^* \geq \delta^{\text{FB}}$. The third item can be proved by following the same steps with a reversed inequality at the contradiction assumption. The first item follows because when $d_0 + \delta^{\text{FB}} = \tilde{d}$, we know

$$R_\nu(\delta^{\text{FB}}, \alpha_\ell) + \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta^{\text{FB}}) - \frac{\kappa}{\rho(d_0 + \delta^{\text{FB}})} \geq R_\nu(\delta, \alpha_\ell) + \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta) - \frac{\kappa}{\rho(d_0 + \delta)}$$

for any $\delta \in [0, \nu]$ and δ^* is the unique optimal solution (see Proposition 5).

Step 2. We next show that $d_0 + \delta^{\text{FB}}$ is an increasing function of d_0 . Using Proposition 2, we obtain the derivative of δ^{FB} with respect to d_0 as follows.

$$\frac{\partial \delta^{\text{FB}}}{\partial d_0} = -\frac{\pi_w''(d_0 + \delta^{\text{FB}})\tau(\alpha_h - \alpha_\ell)}{R_\nu''(\delta^{\text{FB}}, \alpha_\ell) + \pi_w''(d_0 + \delta^{\text{FB}})\tau(\alpha_h - \alpha_\ell)} > -1$$

The inequality here follows from the fact that $R_\nu(\delta, \alpha_\ell)$ and $\pi_w(d_0 + \delta)$ are concave functions. Here the term π_w'' in the numerator is actually the partial derivative of π_w with respect to δ first then with respect to d_0 , and π_w'' in the denominator is the second derivative with respect to δ . Because $d_0 + \delta$ is a linear function, with some abuse of notation, we use π_w'' in both.

If $m(d_{\max}/2) \leq \tilde{d}$, then $\hat{d} = d_{\max}/2$. If $m(d_{\min}) \geq \tilde{d}$, then $\hat{d} = d_{\min}$. If, on the other hand, there exists $\omega \in [d_{\min}, d_{\max}/2]$ such that $m(\omega) = \tilde{d}$, then $\hat{d} = \omega$. Note that the first two cases are degenerate in the sense that there is no d_0 crossing the threshold \hat{d} . Q.E.D.

Proof of Proposition 7. We prove this result using the first-order condition of Problem (6) for $\alpha_1 = \alpha_\ell$ as in the proof of Proposition 5 and the implicit function theorem. Note also that $1/\rho$ is twice-differentiable and with a continuous second-order derivative because π_w and π_s are so and $\pi_w(d) - \pi_s(d) > 0$ for all $d \in [d_{\min}, d_{\max}]$.

To simplify the notation, define $\mu(d_0 + \delta) \triangleq \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta) - \kappa/\rho(d_0 + \delta)$. We use prime over a function (e.g., $\mu'(\cdot)$) to represent the derivative, and double prime to represent second derivative (e.g., $\mu''(\cdot)$). Using this new notation, we represent Problem (6) for $\alpha_1 = \alpha_\ell$ as follows

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \mu(d_0 + \delta_1)$$

and the corresponding first-order condition is

$$\left. \frac{\partial R_\nu(\delta_1, \alpha_\ell)}{\partial \delta_1} \right|_{\delta_1 = \delta^*} + \mu'(d_0 + \delta^*) = 0 \tag{24}$$

Note here that the left-hand side of (24) is a continuously differentiable function of δ^* because $R_\nu(\delta, \alpha_\ell)$ and $\mu(d_0 + \delta)$ are in twice-differentiable with a continuous second-order derivative w.r.t. δ , thus we can use the implicit function theorem. Using the implicit function theorem, we know $\frac{\partial \delta^*}{\partial d_0}$ has the same sign with $\mu''(d_0 + \delta^*)$ because $R_\nu(\delta, \alpha_\ell) + \mu(d_0 + \delta)$ is concave in δ . Since i) $\pi_w(d)$ is concave in d , ii) $1/\rho(d)$ is convex, and iii) $d_0 + \delta$ is a linear function of δ , it follows that $\mu(d_0 + \delta)$ is concave in δ . Therefore, δ^* is nonincreasing in d_0 : $\frac{\partial \delta^*}{\partial d_0} \leq 0$. Q.E.D.

Proof of Proposition 8. When ρ is increasing (decreasing), Theorem 1 implies that $|\delta^* - \delta^{\text{FB}}| = \delta^* - \delta^{\text{FB}}$ ($|\delta^* - \delta^{\text{FB}}| = \delta^{\text{FB}} - \delta^*$) because $\delta^* \geq \delta^{\text{FB}}$ ($\delta^* \leq \delta^{\text{FB}}$). In order to show that the absolute distortion is i) nonincreasing when ρ increasing and ii) nondecreasing when ρ decreasing, we need to prove

$$\frac{\partial \delta^*}{\partial d_0} \leq \frac{\partial \delta^{\text{FB}}}{\partial d_0}. \tag{25}$$

This is because, Propositions 2 and 7, we show that δ^{FB} and δ^* are both nonincreasing in d_0 .

In the proof of this result, we use the fact that $R_\nu(\delta, \alpha_\ell) = \alpha_\ell \frac{\delta}{\nu} \left(1 - \frac{\delta}{\nu}\right)$. Using the implicit function theorem and defining

$$\begin{aligned} i) \text{TERM}_1 &= \pi''_w(d_0 + \delta^{\text{FB}}) \text{ and } \text{TERM}_2 = \frac{-2\alpha_\ell}{\nu^2}, \\ ii) \text{TERM}_a &= \pi''_w(d_0 + \delta^*), \text{ TERM}_b = -\frac{\kappa[\rho'(d_0 + \delta^*)]^2}{[\rho(d_0 + \delta^*)]^3} + \frac{\kappa\rho''(d_0 + \delta^*)}{[\rho(d_0 + \delta^*)]^2} \text{ and } \text{TERM}_c = \frac{-2\alpha_\ell}{\nu^2} \end{aligned}$$

we obtain the following partial derivatives:

$$\frac{\partial \delta^{\text{FB}}}{\partial d_0} = -\frac{\text{TERM}_1}{\text{TERM}_1 + \text{TERM}_2} \text{ and } \frac{\partial \delta^*}{\partial d_0} = -\frac{\text{TERM}_a + \text{TERM}_b}{\text{TERM}_c + \text{TERM}_a + \text{TERM}_b}.$$

Here, we denote by ρ'' the second derivative of ρ . Note that all terms are less than or equal to zero therefore the partial derivatives are so. We next consider the difference of partial derivatives

$$-\frac{\text{TERM}_a + \text{TERM}_b}{\text{TERM}_c + \text{TERM}_a + \text{TERM}_b} + \frac{\text{TERM}_1}{\text{TERM}_1 + \text{TERM}_2} = \frac{\text{TERM}_1 \text{TERM}_c - \text{TERM}_2[\text{TERM}_a + \text{TERM}_b]}{[\text{TERM}_1 + \text{TERM}_2][\text{TERM}_c + \text{TERM}_a + \text{TERM}_b]}. \quad (26)$$

Here, the denominator is nonnegative because it is obtained by multiplying two negative terms. To complete the proof, we need to show that the numerator is nonpositive, i.e., $\text{TERM}_1 \text{TERM}_c \leq \text{TERM}_2[\text{TERM}_a + \text{TERM}_b]$. This inequality is equivalent to

$$\frac{\text{TERM}_1}{\text{TERM}_a + \text{TERM}_b} \leq \frac{\text{TERM}_2}{\text{TERM}_c}.$$

By definition, $\text{TERM}_2 = \text{TERM}_c$ and so $\text{TERM}_2/\text{TERM}_c = 1$.

When $\delta^* \geq \delta^{\text{FB}}$ (ρ increasing) and $\pi''_w(\cdot)$ is nonincreasing, we get $\text{TERM}_1 = \pi''_w(d_0 + \delta^{\text{FB}}) \geq \text{TERM}_a = \pi''_w(d_0 + \delta^*)$. Similarly, $\delta^{\text{FB}} \geq \delta^*$ (ρ decreasing) and $\pi''_w(\cdot)$ is nonincreasing $\text{TERM}_1 = \pi''_w(d_0 + \delta^{\text{FB}}) \geq \text{TERM}_a = \pi''_w(d_0 + \delta^*)$.

$$\text{TERM}_1 \geq \text{TERM}_a + \text{TERM}_b \Rightarrow \frac{\text{TERM}_1}{\text{TERM}_a + \text{TERM}_b} \leq 1 \quad (27)$$

Here, the first inequality follows from the fact that TERM_b is nonpositive, and the second inequality holds because $\text{TERM}_a + \text{TERM}_b \leq 0$. Combining these, we obtain the following inequality and conclude the proof.

$$\frac{\partial \delta^*}{\partial d_0} \leq \frac{\partial \delta^{\text{FB}}}{\partial d_0}.$$

Q.E.D.

Proof of Theorem 2. We initially use the first-order condition of Problem (6) for $\alpha_1 = \alpha_\ell$ and the implicit function theorem to prove this result. The first-order condition is

$$\frac{\alpha_\ell}{\nu} \left(1 - \frac{2\delta^*}{\nu}\right) + \pi''_w(d_0 + \delta^*)(\alpha_h - \alpha_\ell)\tau + \frac{\kappa\rho'(d_0 + \delta^*)}{[\rho(d_0 + \delta^*)]^2} = 0. \quad (28)$$

Because the left-hand side of this equation is decreasing in δ^* (due to concavity), the implicit function theorem implies that the sign of $\frac{\partial \delta^*}{\partial \nu}$ is equal to $\frac{\alpha_\ell}{\nu^2} \left(-1 + \frac{4\delta^*}{\nu}\right)$.

We know that if ρ is increasing $\delta^* \geq \delta^{\text{FB}}$ (see Theorem 1), and $\delta^{\text{FB}} \geq \delta^{\text{M}}$ (see Proposition 1) and $\delta^{\text{M}} = \nu/2$ for the standard uniform distribution (see Lemma 1). Therefore, it follows that $\frac{\partial \delta^*}{\partial \nu} \geq 0$.

If ρ is decreasing, it is possible that δ^* can take values smaller than $\nu/4$. If δ^* is smaller than $\nu/4$ for all $\nu \in (0, d_{\max}/2]$, it follows that δ^* is decreasing in ν therefore $\hat{\nu}$ takes the lowest possible ν value. If δ^* is

larger than $\nu/4$ for all $\nu \in (0, d_{\max}/2]$, it follows that δ^* is increasing in ν therefore $\hat{\nu} = d_{\max}/2$. These two cases are degenerate in the sense that the monotonicity of δ^* does not change in ν .

On the other hand, if δ^* crosses $\nu/4$ at some point $\tilde{\nu} \in (0, d_{\max}/2]$, then δ^* has to be increasing first, i.e., $\delta^* \geq \nu/4$ for all $\nu \in (0, \tilde{\nu}]$, and then has to be decreasing, i.e., $\delta^* \leq \nu/4$ for all $\nu \in [\tilde{\nu}, d_{\max}/2]$. This implies $\hat{\nu} = \tilde{\nu}$. Note that δ^* can cross $\nu/4$ at most once because δ^* is decreasing in ν after crossing but $\nu/4$ is increasing. Q.E.D.

B. Optimal Effort and Price Decisions of the Firm

In this appendix, we first show that if κ is lower than a threshold, then the firm always finds it optimal to retrain the algorithm after observing a low accuracy in the first period. This also implies that retraining the algorithm is profitable in the first-best setting, too. Let $\underline{\Delta\pi} \triangleq \min_{d \in [d_{\min}, d_{\max}]} [\pi_w(d) - \pi_s(d)]$.

PROPOSITION 9. *The firm finds it optimal to retrain the algorithm if $\underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) > \kappa$.*

Proof of Proposition 9. At the time of algorithm retraining decision, the firm has accuracy α_ℓ and the size of the available dataset is d_0 . Incorporating the choice of not retraining, we have the new version of Problem (6) as follows:

$$J_{1b}(\alpha_\ell, d_0) = \max \left[\underbrace{\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + J_{2b}(\alpha_\ell, d_0 + \delta_1)}_{\text{no retraining}}, \underbrace{\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + J_{2a}(\alpha_\ell, d_0 + \delta_1)}_{\text{retraining}} \right] \quad (29)$$

Here, the firm may directly proceed to Period 2b (pricing) with the accuracy on hand, α_ℓ in case of no retraining. The second alternative is collecting data for retraining algorithm as in our model. Using the expected continuation profit J_{2b} from Lemma 1 and J_{2a} from Proposition 4, we rewrite (29) as follows:

$$J_{1b}(\alpha_\ell, d_0) = \max \left[\alpha_\ell \tau, \max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \pi_w(d_0 + \delta_1) \tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d_0 + \delta_1)} \right] + \alpha_\ell \tau$$

We next show that the second term in the square brackets is larger than $\alpha_\ell \tau$ when $\underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) > \kappa$ by finding a uniform positive lower bound for $\pi_w(d)\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d)}$ for all $d \in [d_{\min}, d_{\max}]$. Let $\epsilon \triangleq \underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) - \kappa$, thus $\epsilon > 0$.

$$\begin{aligned} \pi_w(d)\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d)} &= \pi_w(d) \left(\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\pi_w(d) - \pi_s(d)} \right) \\ &\geq \pi_w(d_{\min}) \left(\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\underline{\Delta\pi}} \right) = \frac{\pi_w(d_{\min})\epsilon}{\underline{\Delta\pi}} > 0 \end{aligned}$$

Here, the first equality is obtained using the definition of ρ , the following inequality holds because π_w is increasing and $\underline{\Delta\pi} \leq \pi_w(d) - \pi_s(d)$ for any $d \in [d_{\min}, d_{\max}]$. Using the term $\frac{\pi_w(d_{\min})\epsilon}{\underline{\Delta\pi}}$, we show that

$$\max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \pi_w(d_0 + \delta_1) \tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d_0 + \delta_1)} \geq \max_{\delta_1 \in [0, \nu]} R_\nu(\delta_1, \alpha_\ell) + \frac{\pi_w(d_{\min})\epsilon}{\underline{\Delta\pi}} \geq \alpha_\ell \tau$$

The first inequality follows from the fact that $\frac{\pi_w(d_{\min})\epsilon}{\underline{\Delta\pi}}$ is a uniform lower bound and the second inequality holds because $\frac{\pi_w(d_{\min})\epsilon}{\underline{\Delta\pi}} > 0$ and $\max_{\delta \in [0, \nu]} R_\nu(\delta, \alpha_\ell) = \alpha_\ell \tau$ (see Lemma 1). Q.E.D.

We next show that if the probability of high accuracy π_w and the data impact ρ do not increase or decrease too fast, the firm's optimal price and hence the data size decisions take interior values. Let $\Pi'_{\max} \triangleq \max_{d \in [d_{\min}, d_{\max}]} \pi'_w(d)$ and $\Pi'_{\min} \triangleq \min_{d \in [d_{\min}, d_{\max}]} \pi'_w(d)$ and $D'_{\max} \triangleq \max_{d \in [d_{\min}, d_{\max}]} \rho'(d)$ and $D'_{\min} \triangleq \min_{d \in [d_{\min}, d_{\max}]} \rho'(d)$. Because π_w is concave, in fact $\Pi'_{\max} = \pi'_w(d_{\min})$ and $\Pi'_{\min} = \pi'_w(d_{\max})$.

PROPOSITION 10. *We have that*

- *If $\tau(\alpha_h - \alpha_\ell)\Pi'_{\max} < 2\alpha_\ell/d_{\max}$, then $\delta^{\text{FB}} \in (0, \nu)$ for $\nu \in (0, d_{\max}/2]$.*
- *If $\tau(\alpha_h - \alpha_\ell)\Pi'_{\max} + \kappa D'_{\max}/\underline{\Delta}\pi^2 < 2\alpha_\ell/d_{\max}$ and $\tau(\alpha_h - \alpha_\ell)\Pi'_{\min} + \kappa D'_{\min} > -\alpha_\ell/\nu$, then $\delta^* \in (0, \nu)$ for $\nu \in (0, d_{\max}/2]$.*

Proof of Proposition 10. This result follows from the fact that the derivatives of the objective functions in Problem (12) and Problem (6) for $\alpha_1 = \alpha_\ell$ are positive at $\delta_1 = 0$ and negative at $\delta_1 = \nu$ when the conditions in the statement of the proposition are satisfied. (Recall that the objective functions in both problems are concave in δ_1 , see the proofs of Propositions 1 and 5.)

Note that Proposition 1 implies that $\delta^{\text{FB}} \geq \delta^{\text{M}}$ (and recall that $\delta^{\text{M}} > 0$ from Lemma 1). Thus, we need to check the derivative for the first-best data size only at $\delta_1 = \nu$. Using the definition of $R_\nu(\delta_1, \alpha_\ell)$, we obtain the derivative of the objective function in Problem (12) when $\alpha_1 = \alpha_\ell$ as follows:

$$\alpha_\ell \left(\frac{1}{\nu} - \frac{2\delta_1}{\nu} \right) + \tau(\alpha_h - \alpha_\ell)\pi_w(d_0 + \delta_1). \quad (30)$$

The condition in the first bullet point of the result implies that the term in (30) is negative when evaluated at $\delta_1 = \nu$. Hence, it follows that $\delta^{\text{FB}} \in (0, \nu)$.

Next, we consider δ^* . Similarly, the derivative of the objective function in Problem (6) for $\alpha_1 = \alpha_\ell$ is given as follows:

$$\alpha_\ell \left(\frac{1}{\nu} - \frac{2\delta_1}{\nu} \right) + \tau(\alpha_h - \alpha_\ell)\pi'_w(d_0 + \delta_1) + \frac{\kappa\rho'(d_0 + \delta_1)}{[\rho(d_0 + \delta_1)]^2}. \quad (31)$$

The first condition in the second bullet point guarantees that the term in (31) evaluated at $\delta_1 = \nu$ is negative for any $\nu \in (0, d_{\max}/2]$ because $\pi'_w(d_0 + \nu) \leq \Pi'_{\max}$ and $\rho(d_0 + \delta) \geq \underline{\Delta}\pi$. The second condition in the same bullet point guarantees that the term in (31) is positive when evaluated at $\delta_1 = 0$. Therefore, it follows that $\delta^* \in (0, \nu)$. Q.E.D.

Although the conditions in Proposition 10 are sufficient to guarantee that the optimal solutions δ^* and δ^{FB} are interior, they are in fact loose because we do not assume any parametric form of functions π_w and ρ . For given parametric forms of functions π_w and ρ , these conditions can be improved. Alternatively, it is straightforward to check if the optimal solutions δ^* and δ^{FB} are interior for any given set of parameters as in the numerical examples provided in the main body of our paper.

Recent ESMT Working Papers

	ESMT No.
Beyond retail stores: Managing product proliferation along the supply chain	19-02 (R1)
Işık Biçer, Schulich School of Business, York University Florian Lücker, Cass Business School, City, University of London Tamer Boyacı, ESMT Berlin	
The impact of EU cartel policy reforms on the timing of settlements in private follow-on damages disputes: An empirical assessment of cases from 2001 to 2015	19-03
Hans W. Friederiszick, ESMT Berlin and E.CA Economics Linda Gratz, E.CA Economics Michael Rauber, E.CA Economics	
Marginality, dividends, and the value in games with externalities	19-01
Frank Huettner, ESMT Berlin André Casajus, HHL Leipzig Graduate School of Management	
Consumer choice under limited attention when alternatives have different information costs	16-04 (R3)
Frank Huettner, ESMT Berlin Tamer Boyacı, ESMT Berlin Yalçın Akçay, Melbourne Business School	
Opaque queues: Service systems with rationally inattentive customers	18-04
Caner Canyakmaz, ESMT Berlin Tamer Boyacı, ESMT Berlin	