

October 19, 2020

ESMT Working Paper 20-01 (R1)

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

Huseyin Gurkan, ESMT Berlin

Francis de Véricourt, ESMT Berlin

Revised version

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 (<https://esmt.berlin/>) 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

Huseyin Gurkan

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

Francis de Véricourt

ESMT Berlin, Germany, francis.devericourt@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. Interestingly, these distortions sometimes improve social welfare, which accounts for the customer surplus and profits of both the firm and provider. Further, the interaction between incentive issues and the positive externalities of the AI Flywheel effect have important implications for the firm's data collection strategy. In particular, the firm can boost its profit by increasing the product's capacity to acquire usage data only up to a certain level. If the product collects too much data per user, the firm's profit may actually decrease. As a result, the firm should consider reducing its product's data acquisition capacity when its initial dataset to train the algorithm is large enough.

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

History: October 19, 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 (Sarıkaya 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 the startup Blue River Technology famously established their first dataset manually to train an AI system that would distinguish weeds from crops, a crucial step to optimize pesticide sprays 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 the accuracy of algorithms 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 startups and small organizations (Nicolaus Henke et al. 2016). This shortage has given rise to a thriving outsourcing industry (Research Nester Pvt. Ltd 2019), and many startups have achieved success by outsourcing their technology (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 may expose the firm to threats (Kendra et al. 2019, Bursztein 2018). If these algorithms are not developed with care, they only learn 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 if the firm does not outsource the algorithm but instead employs an expert. Indeed, the details of ML algorithms and their outputs notoriously suffer 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, that is, 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. These effects also affect social welfare and customer surplus, which improve (resp. deteriorate) when the impact of the provider’s effort on accuracy increases (resp. decreases) with more data.

More importantly, the interaction between the positive externalities of the AI Flywheel effect and the incentive issues can alter the firm’s data collection strategy. In particular, a key and specific design choice for data products is the amount of data that the product can collect on its user, which we refer to as the data acquisition capacity in the following (Spencer 1990, Kos et al. 2012). The firm can increase this amount by relying on third-party services in the case of mobile and web applications (Deshpande 2019) or increasing the capacity of embedded sensors in the case of physical products (McGrath and Scanail 2013). This, in turn, should provide the firm with more data overall to improve its algorithm and ultimately increase profits.

We find, however, that increasing the product’s data acquisition capacity sometimes reduces the total amount of data collected by the firm. This happens when the provider’s effort has a less significant impact on accuracy for larger volumes of data. In this case, increasing the data acquisition capacity does boost profit, but up to a certain point. If the product collects too much data per user, the profit may actually decrease. In other word, the data acquisition capacity has a unimodal effect on the firm’s profit. This further implies that the firm should consider reducing the product’s data acquisition capacity once its initial dataset to train an algorithm is large enough.

Taken together, these results characterize how the interaction between the incentive issues and the positive adoption externalities of the AI Flywheel effect significantly affects the firm’s pricing and data collection strategies. One 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 acquisition 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 lower revenues. Further, the initial data size and the acquisition capacity together characterize the potential strength of the AI Flywheel effect. The initial data size specifies the firm's starting point in the virtuous cycle, while the acquisition capacity 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 first characterize the optimal contract and equilibrium payments to the provider in the presence of incentive issues (see Section 3). We then briefly study the optimal decisions of the firm at first best, that is, in the presence of the AI Flywheel effect but in the absence of incentive issues (see Section 4.1). Under the 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 4.2). The impact on the social welfare and customer surplus exhibits similar distortions. In particular, social welfare and customer surplus actually improve in the presence of incentive issues when the data impact is increasing (see Section 4.3).

Finally, we show that improving the product’s data acquisition capacity increases both the total amount of collected data and the firm’s profit. In contrast, improving the acquisition capacity has a unimodal effect on the collected data and profit when the data impact decreases (see Section 5). In this case, a threshold exists such that increasing the product’s data acquisition capacity boosts the firm’s profit if and only if the initial size of the dataset is below this threshold. We discuss the managerial implications of these results and future research directions in the conclusion (see Section 6).

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, that is, the enhancement of quality. In addition, we consider moral hazard issues, which is not the focus of this stream of research.

Another related stream of literature studies the positive network effects (or positive adoption externalities) that customers derive from the consumption of the same product by others (see Katz and Shapiro 1994 and Shy 2011 for a general review). In particular, this literature considers the dynamic monopoly pricing of durable goods with network effects that are close to our first-best benchmark (see for instance Cabral et al. 1999, Bensaid and Lesne 1996, Gabszewicz and Garcia 2008). In these studies, customers are strategic and long-lived, and thus have incentives to postpone their purchases in order to benefit from the network effects generated by others' consumptions. This induces the firm to decrease the price in the first period.

Our work differs from this stream of research for two reasons. First, we explore whether a firm should strengthen the network effects, which corresponds to improving the data acquisition capacity in our setup. Although the acquisition capacity is a natural design choice in our setting, firms cannot easily affect the intensity of the network effects in the aforementioned literature. Second, these studies are not concerned with incentive issues on the provider side. In our setup, however,

the amount of collected data interacts with both the positive adoption externalities of the AI Flywheel effect and the incentive issues that developing the product brings about.

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, 2013) explores similar settings, where the probabilities of success are assumed to change across periods.

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, 2013). By contrast, these probabilities are fully endogenous in our setup.

In this stream of research, Dai and Jerath (2013, 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 discussed 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 Dai and Jerath (2013), 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. More generally, deviations from the first-best setting are observed in other contexts as well, (see, e.g., Alles et al. 1995 for over/understocking decisions in just-in-time production systems). A common theme of the moral hazard literature is that it is attractive for the principal to take actions that mitigate incentive issues, and the unobservability of the agent's effort in particular.

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. Moreover, ours is the first to articulate this theme in the context of data analytic. Specifically, the AI Flywheel effect endogenously interacts with incentive issues and both are affected by data. Thus, the firm’s ability to leverage the AI Flywheel effect is constrained by the incentive misalignment between the firm and the provider.

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) at 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 $\alpha_h > \alpha_\ell$.¹ Given

¹ Accuracy in our setup designates the overall accuracy, which includes both sensitivity and specificity, and this metrics is typically used to measure the performance of ML algorithms (see, e.g., Igual and Seguí 2017).

the 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 the 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 .

A key attribute of data-enhanced products is the amount of data that the product can collect on its user. The firm may increase this amount by relying on third-party services in the case of mobile and web applications (Deshpande 2019) or increasing the capacity of embedded sensors in the case of physical products (McGrath and Scanail 2013). We refer to this quantity as the data acquisition capacity $\nu > 0$ in the following, which corresponds to the expected amount of data the firm collects 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 acquisition capacity 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 acquisition capacity 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 $\pi_e(d_0)$.

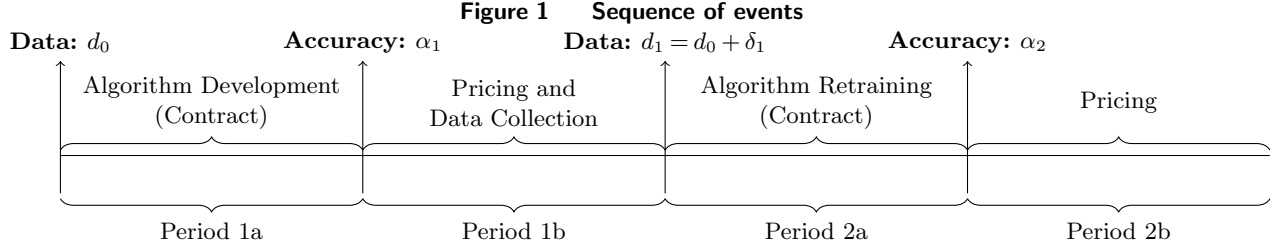
Period 1b – Pricing and Data Collection. Given accuracy α_1 , the firm prices and markets the product. The firm collects δ_1 such that the total size of the dataset becomes $d_1 = d_0 + \delta_1$ and generates revenue $R_\nu(\delta_1, \alpha_1)$.

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.)² 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, that is, $\pi_e(d_1) > \pi_e(d_0)$, 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

² Our model and results easily extend to the case where accuracy is cumulative, that is, in situations where accuracy can be further improved when $\alpha_1 = \alpha_h$.



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 $d_{\min} \leq d_0 \leq d_{\max}$ and $0 < \nu \leq d_{\max} - d_0$.

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, that is,

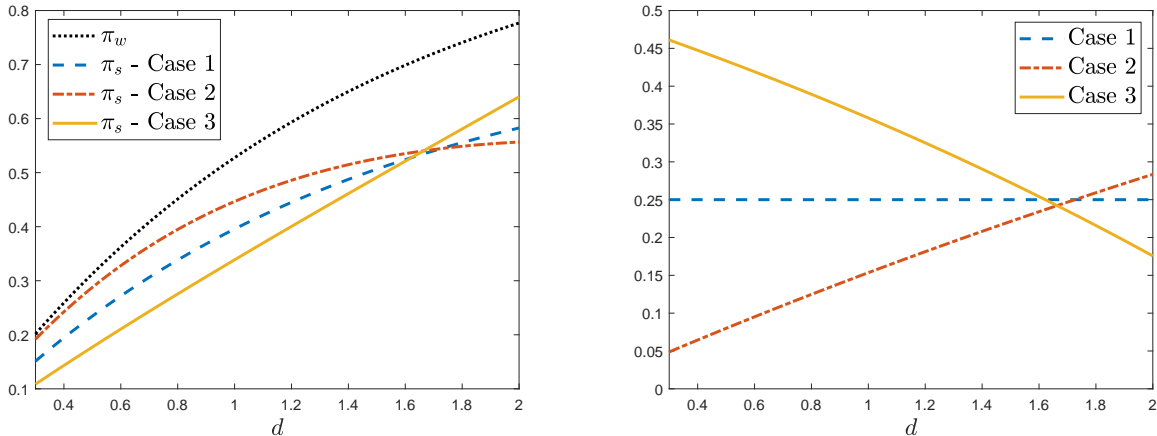
$$\rho(d) \triangleq \frac{\pi_w(d) - \pi_s(d)}{\pi_w(d)}. \quad (1)$$

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.

Many corroborative evidences support the assumption that the data impact is decreasing in a variety of practical settings. Specifically, empirical evidence demonstrates that the difference in accuracy between sophisticated ML algorithms (which require more effort to develop and train), and more crude ones decrease with the amount of data on which they are trained (Banko and Brill

Figure 2 The different regimes of the data impact



(a) Effect of d on π_w and π_s

(b) Effect of d on ρ

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).

2001; Halevy et al. 2009). Practitioners have also recognized this phenomena, see, for example, Gutierrez (2016).³

Yet, certain circumstances may justify assuming and increasing data impact. This is the case, for instance, when biases in the training data sets are of particular concern. The impact of these biases are typically more pronounced for larger datasets.⁴ In this case, the effort to improve the algorithm has a bigger impact with more data. In addition, the performance of deep learning algorithms seem to outperform very basic statistical models (such as logistic regressions) more when datasets are larger (see Hackathorn 2018, for instance). Thus, the data impact should increase when the firm does not have the skills to run nor check these more basic methods, and hence cannot detect this form of shirking.

³ For instance, the following quote is from a data scientist at The Trade Desk “The bigger the training set, the better the model. But less obviously, and more importantly, the difference between a fancy algorithm and a simple one decreases with more data.” Gutierrez (2016).

⁴ A well-known example is Microsoft’s AI-based chatbot in Twitter that caused unintended tweets as the size of the dataset feeding the algorithm increased, see Vincent (2016).

Figure 2 depicts different examples corresponding to these different regimes. In all examples, probability $\pi_w(\cdot)$ (the black dotted 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).

Finally, 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 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 and has commitment power across periods. Thus, the firm needs to offer a single contract which determines the provider's payments in the first period along with the expected payoff in the second period through the total amount of collected data δ_1 and the second period payments (we present the noncommitment case in Appendix D).

Accuracy realizations α_t , $t = 1, 2$, are contractable but effort is not. In particular, the product may track (ex-post) accuracy once it is distributed, or the firm can always hold some data to test the accuracy. Alternatively, the firm may contract on revenues, which is equivalent to contracting on accuracy in our setup.

In contrast, ML algorithms are notoriously difficult to explain to non-experts (Lipton 2016, Ribeiro et al. 2016), rendering the expert's efforts to improve accuracy difficult to observe and

contract on (Gromb and Martimort 2007). Thus, the contract's payments in each period depend on the public history and are contingent upon the accuracy realizations.⁵ Specifically, the history in Period 1a decreases 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.⁶ 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). We then 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

⁵ In our setup, the accuracy realization affects the demand per Section 2. Thus, the contract's payments can equivalently be contingent upon the demand.

⁶ Recall that these last payments are only meaningful when $\alpha_1 = \alpha_\ell$, as no contract is required in the second period when $\alpha_1 = \alpha_h$.

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 expected continuation profit in Period 2a, given payments $x_{2h}, x_{2\ell}$, and accuracy α_1 and data size d_1 , such that

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

$$J_{2a}(\alpha_h, d_1) = J_{2b}(\alpha_h, d_1). \quad (4)$$

When $\alpha_1 = \alpha_\ell$, the firm needs to set payments such that the provider has enough incentives to exert effort, as will be formalized by incentive constraint (9). These payments are then deduced 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.

Moving to the first period, we denote by $J_{1b}(\alpha_1, d_0, \delta_1)$ the firm's expected continuation profit in Period 1b given accuracy α_1 , data size d_0 , and collected data δ_1 , such that

$$J_{1b}(\alpha_1, d_0, \delta_1) = R_\nu(\delta_1, \alpha_1) + J_{2a}(\alpha_1, d_0 + \delta_1) \text{ for } \alpha_1 \in \{\alpha_h, \alpha_\ell\}. \quad (5)$$

When $\alpha_1 = \alpha_\ell$, the choice of data δ_1 (or equivalently the price) affects current revenues directly and future ones indirectly by increasing the dataset size to $d_0 + \delta_1$.

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

$$J_{1a}(d_0) = \max_{\substack{x_{1h}, x_{1\ell} \geq 0 \\ x_{2h}, x_{2\ell} \geq 0 \\ \delta_{1h}, \delta_{1\ell} \in [0, \nu] \\ \text{s.t.}}} \pi_w(d_0)[J_{1b}(\alpha_h, d_0, \delta_{1h}) - x_{1h}] + \bar{\pi}_w(d_0)[J_{1b}(\alpha_\ell, d_0, \delta_{1\ell}) - x_{1\ell}] \quad (6)$$

$$\pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[x_{1\ell} + J_p(d_0)] - \kappa \geq \pi_s(d_0)x_{1h} + \bar{\pi}_s(d_0)[x_{1\ell} + J_p(d_0)] \quad (7)$$

$$J_p(d_0) = \pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} - \kappa. \quad (8)$$

$$\pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} - \kappa \geq \pi_s(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_s(d_0 + \delta_{1\ell})x_{2\ell} \quad (9)$$

The firm needs to provide incentive to the provider for both periods. Therefore, 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 (7) 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 (8). A key aspect of our setup is that the distribution of these future payments, $\pi_w(d_1)$, depends on the choice of $\delta_{1\ell}$ since $d_1 = d_0 + \delta_{1\ell}$. Thus, the choice of $\delta_{1\ell}$ (or equivalently price) not only makes the tradeoff between present and future revenues as in (5) but also determines the intensity of the incentive issue. 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 3.

3 Optimal Decisions of the Firm

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 $\delta_{1\ell}, \delta_{1h}$ (and the corresponding price $p_{1\ell}, p_{1h}$) as well as all payments $x_{1\ell}, x_{1h}, x_{2\ell}$, and x_{2h} in J_{1a} .

The optimization problem of J_{1a} in (6) corresponds to maximizing the expected profit, that is, the total revenue net of the payments motivating the provider to exert effort. This problem can be solved in three steps. Starting from Period 2b, we initially consider the myopic problem J_{2b} whose solution is characterized in Lemma 1. Next, fixing the amount of collected data $\delta_{1\ell}$ and δ_{1h} , we characterize the corresponding payments. The following lemma (Lemma 2) provides the optimal payments which motivate to exert effort for given $\delta_{1\ell}, \delta_{1h}$. Finally, we characterize the optimal amount of collected data which maximizes the total expected profit of the firm.

Specifically, the firm's problem in Period 2b corresponds to the myopic problem in (2). Straight-forward calculations then lead to 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.

LEMMA 2. *Fix $\delta_{1\ell}$ and δ_{1h} in (6), then the corresponding optimal payments are $x_{1\ell}(\delta_{1\ell}) = x_{2\ell}(\delta_{1\ell}) = 0$, and*

$$x_{1h}(\delta_{1\ell}) = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + \left(\frac{\kappa}{\rho(d_0 + \delta_{1\ell})} - \kappa \right) \text{ and } x_{2h}(\delta_{1\ell}) = \frac{\kappa}{\pi_w(d_0 + \delta_{1\ell}) - \pi_s(d_0 + \delta_{1\ell})}.$$

The firm's total expected payment to the provider is then given by

$$\frac{\kappa}{\rho(d_0)} + \frac{\kappa}{\rho(d_0 + \delta_{1\ell})}.$$

In the first period, the firm needs to account for the provider's future expected profits contrary to the payments in Period 2. Specifically, optimal payment $x_{1h}(\delta_{1\ell})$ 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 $x_{1h}(\delta_{1\ell})$ captures the cost due to the scarcity of AI service providers in the market. In both periods, the ex-post payments to the provider depend on the amount of data $\delta_{1\ell}$. Because there is no need for further improvement after a high accuracy α_h in the first period, payments depend only on $\delta_{1\ell}$.

MLRP ensures that under the optimal contract, realizations of higher value (α_h in our setup) are more rewarded ($x_{2h} > x_{2l}$). Thus, MLRP concerns the ex-post payments to the provider. In contrast, Lemma 2 shows that data impact $\rho(\cdot)$ determines the ex-ante provider's rent $\kappa[1/\rho(d_0) + 1/\rho(d_0 + \delta_{1\ell})]$.

Using Lemma 1 and 2, Problem (6) can be reduced to choosing data sizes $\delta_{1\ell}$ and δ_{1h} . We next show that the optimal data sizes solving Problem (6) and hence the corresponding optimal payments are unique.

PROPOSITION 1. *The optimal solution to Problem (6) is unique such that $\delta_{1h} = \delta^M$ and $\delta_{1\ell} = \delta^*$ where δ^* is the unique optimal solution of*

$$\max_{\delta \in [0, \nu]} \bar{\pi}_w(d_0) [R_\nu(\delta, \alpha_\ell) + \pi_w(d_0 + \delta)(\alpha_h - \alpha_\ell)\tau] - \frac{\kappa}{\rho(d_0 + \delta)}.$$

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 myopic price. We refer to δ^* as the optimal solution $\delta_{1\ell}$ of Problem (6), and to p^* as the price that yields data size δ^* (see Section 2).

This proposition further reveals how the main effects captured by our model endogenously determine data collection size δ^* . Indeed, deviating δ away from δ^M decreases the firm's current period revenue $R_\nu(\delta, \alpha_\ell)$. But, increasing δ always boosts the positive externalities of the AI Flywheel effect ($\pi_w(d_0 + \delta)(\alpha_h - \alpha_\ell)\tau$). Moreover, additional data may exacerbate or mitigate the incentive issues and associated agency cost $\kappa/\rho(d_0 + \delta)$ depending on the monotonicity of the data impact. The choices of δ^* essentially balances these different tradeoffs simultaneously.

4 Impact of Incentive Issues

We first analyze the impact of incentive issues on the optimal decisions of the firm along with the social welfare and consumer surplus for a fixed pair of d_0 and ν . To do so, we introduce a comparison benchmark, the first-best setting, where the firm does not face any incentive issues.

4.1 First-Best Benchmark

In the first-best setting, the firm has the capability to develop the algorithm and does not face any incentive issues. The first-best problem then corresponds to Problem (6) without incentive constraints (7) and (9), but where the firm directly incurs cost κ . Specifically, the problem without incentive issues, denoted by $J_{1a}^{FB}(d_0)$, is obtained by setting payments to κ (i.e., $x_{1h} = x_{1\ell} = x_{2h} = x_{2\ell} = \kappa$) and removing incentive constraints (7) and (9) in Problem (6). The following proposition characterizes the firm's optimal decision at first-best.

PROPOSITION 2. *The optimal solution to $J_{1a}^{FB}(d_0)$ is unique such that $\delta_{1h} = \delta^M$ and $\delta_{1\ell} = \delta^{FB}$ where δ^{FB} is the unique solution $\max_{\delta \in [0, \nu]} R_\nu(\delta, \alpha_\ell) + \pi_w(d_0 + \delta)(\alpha_h - \alpha_\ell)\tau$, and $\delta^{FB} > \delta^M$.*

We denote by δ^{FB} the optimal value of $\delta_{1\ell}$ maximizing $J_{1a}^{\text{FB}}(d_0)$. We also refer to p^{FB} as the corresponding price that yields data size δ^{FB} . If the firm succeeds in developing the initial algorithm with 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, that is, $\delta^{\text{FB}} > \delta^{\text{M}}$, increasing the probability of high accuracy and hence expected profit in the next period.

To the extent that we consider a dynamic model, these distortions away from the myopic decision ($\delta^{\text{FB}} > \delta^{\text{M}}$) are to be expected. They occur in many different settings, such as in inventory problems (see Heese and Swaminathan 2010, for example), but are especially present in problems of pricing under positive adoption externalities (see Section 1.1).

Finally, the firm's pricing and data collection decisions p^{FB} and δ^{FB} determine the social welfare and customer surplus at first best, which we denote by W^{FB} and C^{FB} , respectively. Customer surplus C^{FB} is the total valuations net of the product price of all customers who make a purchase. Social welfare W^{FB} is then the sum of customer surplus C^{FB} and the firm's profit net of effort cost κ .

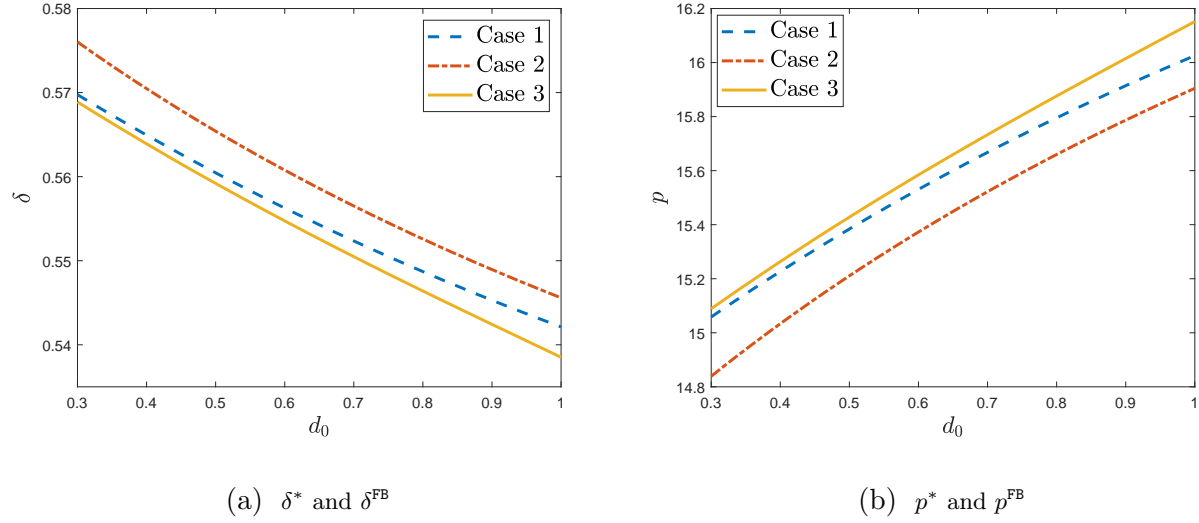
4.2 Overpricing, Underpricing, and Optimal Data Collection

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

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, $\delta^{FB} = \delta^*$ and $p^{FB} = p^*$ are in Figure 3a and Figure 3b, respectively.

required. By contrast, if this impact intensifies with more data (ρ is increasing), the firm needs to underprice $p^* \leq p^{FB}$ in order to collect more data $\delta^* \geq \delta^{FB}$. Finally, the firm needs to overprice $p^* \geq p^{FB}$ 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.

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 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 compared to first-best. This situation happens, for instance, when $p^{\text{M}} > p^* > p^{\text{FB}}$. Overall, the incentive issues prevent the firm from fully leveraging the AI Flywheel effect in this regime.

4.3 Social Welfare and Customer Surplus

Theorem 1 uncovers effects that have further implications for social welfare and customer surplus in the presence of incentive issues, which we denote by W^* and C^* , respectively. In this case, the social surplus corresponds to the sum of the firm's profit, customer surplus C^* , and the provider's profit.

The next result shows that the social benefits are monotone in the firm's data collection decision at equilibrium.

PROPOSITION 3. *Social welfare W^* and consumer surplus C^* are nondecreasing in δ^* . Furthermore, $W^* = W^{\text{FB}}$ and $C^* = C^{\text{FB}}$ for any $\delta^* = \delta^{\text{FB}}$.*

That a higher δ^* should improve the social benefits at the equilibrium is not obvious a priori. In particular, a larger data collection corresponds to a higher expected accuracy in the next period.

This, however, also increases the expected equilibrium price in Period 2b (see Lemma 1). Our result shows that the latter effect is dominated by the former and thus both the social welfare and the customer surplus increase with data collection δ^* .

Thus, collecting more data at the equilibrium boosts the social benefits. Based on Theorem 1, we can then characterize how the incentive issues affect the social welfare and customer surplus. In fact, the following result indicates that incentive issues may *increase* these benefits.

PROPOSITION 4. *We have*

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

In other words, the firm's distortion from the first-best benchmark to mitigate the effects of incentive issues improves the social welfare and the customer surplus when the data impact is increasing. In contrast, this distortion hurts the social welfare when the data impact is decreasing.

The effect of incentive issues have further implications for these societal benefits. Specifically, our setup considers a unique provider, with which the firm contracts in both period. It is not clear, however, how the previous societal benefits change when the firm can contract with different providers in different periods. (Our results also hold overall for the competitive provider market setting, see Appendix C.) The next result shows that a more competitive market of providers may actually hurt the social benefits.

PROPOSITION 5. *Let \tilde{W}^* and \tilde{C}^* be the social welfare and, respectively the customer surplus in a setting where the firm changes providers in each period.*

1. *If ρ is increasing, $W^* \geq \tilde{W}^*$ and $C^* \geq \tilde{C}^*$.*
2. *If ρ is decreasing, $W^* \leq \tilde{W}^*$ and $C^* \leq \tilde{C}^*$.*

Thus, a more competitive provider market hurts societal benefits more if the data impact increases. This is because the firm suffers more from the moral hazard of the provider in the absence of competition. Therefore, the firm collects more data when the data impact is increasing to mitigate incentive issues compared to the competitive provider market.

5 Impact of the Data Acquisition Capacity

One key product attribute that firms commonly seek to improve is the amount of data that the product collects on its user. In our setup, this corresponds to increasing data acquisition capacity ν . In the following, we explore how this improvement affects the firm's data collection strategy and expected profit. The main finding of this analysis is that data acquisition capacity ν can have a non-monotonic impact on the firm's profit.

5.1 Optimal Data Collection

Without incentive issues, a higher data acquisition capacity induces the firm to collect even more data under the AI Flywheel effect, as seen in the next result.

PROPOSITION 6. *The first-best data size δ^{FB} is nondecreasing in data acquisition capacity ν .*

If the firm considers designing a product that acquires more data per user, the firm collects in aggregate even more data (δ^{FB} is nondecreasing in ν). 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.

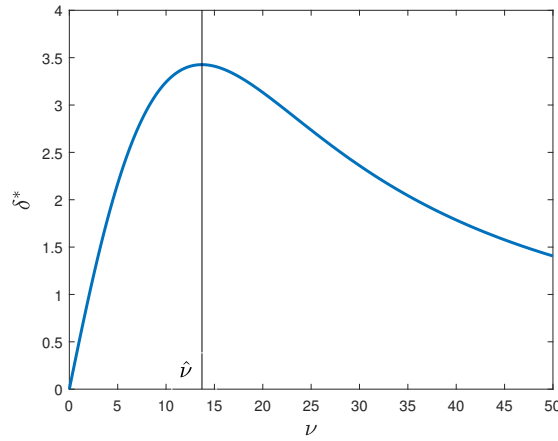
The need to mitigate incentive issues, however, sometimes reverses this effect and pushes the firm to collect *less* data when the product's data acquisition capacity is higher. The next theorem, one of our main results, formalizes this finding.

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.*

Thus, increasing the data acquisition capacity induces the firm to collect more data overall (δ^* is increasing), when the data impact is also increasing. This effect, however, may continue to hold even when the data impact is decreasing, that is, when more data actually exacerbates the incentive issues. This is the case as long as the data acquisition capacity is not too large ($\nu < \hat{\nu}$). When the

Figure 4 **Effect of the product's data acquisition capacity on the firm's data collection decision when more data exacerbate the incentive issues**



Note. Probabilities generating this example are $\pi_w(d) = 1 - 1/\sqrt{d}$, $\pi_s(d) = \pi_w(d)(1 - 1/(d+1))$ and the data impact is $\rho(d) = 1/(1+d)$. The size of the initial dataset is $d_0 = 10$.

data acquisition capacity becomes large ($\nu \geq \hat{\nu}$), acquiring more data per user actually induces the firm to collect less data overall (δ^* is decreasing).

Figure 4 illustrates this point and depicts the effect of ν on δ^* when the data impact is decreasing. The figure shows nonmonotonic behavior in the acquisition capacity ν . In particular, this contrasts with the first-best benchmark δ^{FB} , which always increases in ν according to Proposition 6. Threshold $\hat{\nu}$ corresponds to the straight black line in Figure 4), with $\hat{\nu} \cong 13.7$.

More specifically, when the acquisition capacity 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 6). 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 acquisition capacity is small, the amount of collected data $\delta = \nu q$ remains small for any quantity q , which yields low agency costs. As the acquisition capacity increases, boosting the AI Flywheel effect dominates the increases in agency costs, and data size δ^* increases as a result.

When the acquisition capacity 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^* .

5.2 The Firm's Profit

To the extent that acquisition capacity ν determines the firm's data collection strategy per Theorem 2, improving this capability ultimately affects the firm's profit. The next result – one of the main findings of the paper – characterizes this effect.

THEOREM 3. *We have:*

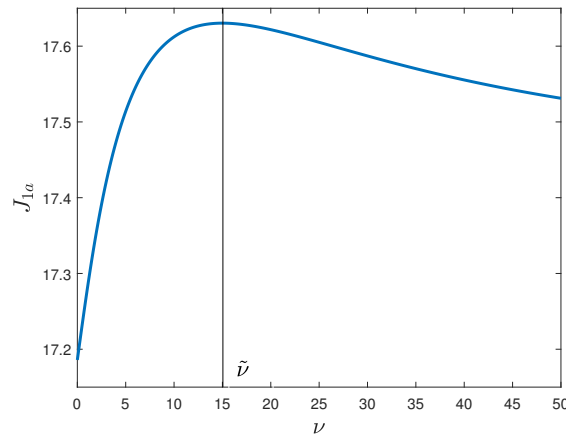
1. *If $\rho(\cdot)$ is increasing, the expected profit of the firm is nondecreasing in ν .*
2. *If $\rho(\cdot)$ is decreasing, a unique threshold $\tilde{\nu}$ exists such that the expected profit is nondecreasing in ν if $\nu \geq \tilde{\nu}$, and is nonincreasing otherwise.*

Thus, when the data impact is increasing, the firm's profit is higher when the product can acquire more data on its user. In fact, this effect continues to hold even when more data exacerbates the incentive issues, as long as the data acquisition capacity is not too large ($\nu < \tilde{\nu}$). However, if the data acquisition capacity is large enough ($\nu \geq \tilde{\nu}$), boosting the product's capacity to collect data actually hurts the firm's profit.

Three main components of our model jointly shape the impact of the data acquisition capacity on the firm's profit: the AI Flywheel effect, pricing for revenue, and incentive issues. To leverage the AI Flywheel effect more, the firm needs to collect more data. Therefore, optimal data collection δ^* deviates from myopic data collection δ^M , and optimal price p^* from myopic price p^M (see Section 4.2), which hurts revenues in the first period. Improving the data acquisition capacity ν makes it possible for the firm to push the price toward the myopic one, thus boosting revenues in the current period without affecting the amount of data collected to retrain the algorithm in the next period. This positive effect continues to hold for larger values of the acquisition capacity when the data impact is increasing (per the first item of Theorem 3).

When the data impact is decreasing, however, the previous positive effect holds until the acquisition capacity hits threshold $\tilde{\nu}$, which is exactly the value at which the firm charges the myopic price ($p^* = p^M$). In this case, the firm has maximized the revenue it can collect from pricing in the first period. Increasing the data acquisition capacity further exacerbates the incentive issues in the second one. In addition, the marginal contribution to the accuracy of the additional collected data is decreasing (recall that π_w is increasing but concave). Thus, exceeding $\tilde{\nu}$ generates no additional revenue in the current period, while the improvements in accuracy and hence revenues in the second one are offset by the increase in agency costs. The firm's profit decreases as a result. Figure 5

Figure 5 Effect of the product's data acquisition capacity on the firm's profit when more data exacerbate the incentive issues



Note. Probabilities generating this example are $\pi_w(d) = 1 - 1/\sqrt{d}$, $\pi_s(d) = \pi_w(d)(1 - 1/(d+1))$ and the data impact is $\rho(d) = 1/(1+d)$, which is decreasing. The size of the initial dataset is $d_0 = 5$.

illustrates this last point and depicts the effect of ν on the firm's expected profit J_{1a} when the data impact is decreasing. The figure shows that the data acquisition capacity ν has a unimodal effect on profit. Threshold $\hat{\nu}$ corresponds to the straight black line in Figure 5, with $\tilde{\nu} \approx 15$.

This unimodal relationship between the data acquisition capacity and the firm's profit has further implications for the role of initial data size d_0 , as presented in the following result.

PROPOSITION 7. *If $\rho(\cdot)$ is decreasing, there exists a threshold \tilde{d}_0 such that the expected profit of the firm is decreasing in ν when $d_0 \geq \tilde{d}_0$, and is unimodal otherwise.*

Thus, when the initial dataset of the firm is large enough ($d_0 \geq \tilde{d}_0$), increasing the product's acquisition capacity always hurts profits when the data impact is decreasing. Because the initial dataset is large, the firm needs to rely less on the AI Flywheel effect to collect data for the second period. The firm nonetheless charges a price above the myopic one to mitigate the incentive issues. If the data acquisition capacity improves, the incentive issues intensify and the firm increases the price away from the myopic one, which decreases revenues.

In contrast, when the initial size of the dataset is not too large ($d_0 < \tilde{d}_0$), the relationship between the acquisition capacity and profit is unimodal. This is because when initial size d_0 and acquisition capacity ν are small, the firm needs to leverage the AI Flywheel effect to collect more data and thus prices the product below the myopic price. As acquisition capacity ν increases, the firm can then increase the optimal price closer to the myopic revenues while collecting the same amount of data for the next period. The total expected profit increases as a result, until the optimal price reaches the myopic one. In this case, we retrieve the previous effects and profits start to decrease in the acquisition rate.

6 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 the 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 regulate the impact of the data acquisition capacity on the firm's profit and decisions. In particular, when the data impact decreases, we find that improving the capacity to acquire user data may hurt profits

and induce the firm to actually collect less data overall. These regimes have further implications for society. Notably, incentive issues improve social welfare when the data impact is increasing.

Improving a product’s capacity of acquiring usage data is a key but costly and at times challenging design choice in practice (e.g., Boonstra et al. 2018). As such, our insights on when a firm should seek to improve the acquisition capacity address an important managerial problem. In essence, the acquisition capacity determines the strength of the positive adoption externalities of the AI Flywheel effect in our setup. In the context of classical products, this corresponds to the intensity of the customer social network or how fast information circulates in the market (Shy 2011), which the firm cannot easily influence. Nonetheless, the stronger the network effect in classical adoption settings, the better off the firm typically is (Gabszewicz and Garcia 2008). To some extent, this remains true for data products. In our setup, entrepreneurs who seek to leverage the AI Flywheel effect should indeed consider enhancing the data collection capability of their products, but only up to a certain level. Allowing a data product to collect too much usage data may sometimes damage the firm’s profits.

Our 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, which hurts their profits (per Theorem 2 and Theorem 3, respectively).

From a technical perspective, our model can be extended in a variety of ways. For instance, exploring how information asymmetry regarding the provider’s skill interacts with the AI Flywheel effect constitutes a fruitful future research direction. Similarly, future research could explore our setup with a longer time horizon.

More generally, we believe that our work opens up more fundamental 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.

Overall, our results provide the first insights on how firms can leverage the AI Flywheel effect, to the best of our knowledge. 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, Makhdoumi 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* .
- Alles M, Datar SM, Lambert RA (1995) Moral hazard and management control in just-in-time settings. *Journal of Accounting Research* 33:177–204.
- 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).
- Bensaid B, Lesne JP (1996) Dynamic monopoly pricing with network externalities. *International Journal of Industrial Organization* 14(6):837–855.
- Bimpikis K, Crampton D, Tahbaz-Salehi A (2019) Information sale and competition. *Management Science* 65(6):2646–2664.
- Boonstra TW, Nicholas J, Wong QJ, Shaw F, Townsend S, Christensen H (2018) Using mobile phone sensor technology for mental health research: integrated analysis to identify hidden challenges and potential solutions. *Journal of medical Internet research* 20(7):e10131.

- Bursztein E (2018) Attacks against machine learning ? an overview. elie.net/blog/ai/attacks-against-machine-learning-an-overview/ [Online; accessed 27-January-2020].
- Cabral LM, Salant DJ, Woroch GA (1999) Monopoly pricing with network externalities. *International Journal of Industrial Organization* 17(2):199–214.
- 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 (2013) Salesforce compensation with inventory considerations. *Management Science* 59(11):2490–2501.
- 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.
- Gabszewicz JJ, Garcia F (2008) A note on expanding networks and monopoly pricing. *Economics Letters* 98(1):9–15.
- 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].
- Gromb D, Martimort D (2007) Collusion and the organization of delegated expertise. *Journal of economic theory* 137(1):271–299.
- Gutierrez D (2016) Data-first machine learning. <http://www.inetservicescloud.com/data-first-machine-learning/> [Online; accessed 11-September-2020].

- Hackathorn R (2018) How managers should prepare for deep learning: New paradigms. <https://towardsdatascience.com/how-managers-should-prepare-for-deep-learning-new-paradigms-28de63054ea6> [Online; accessed 15-09-2020].
- 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>.
- Heese HS, Swaminathan JM (2010) Inventory and sales effort management under unobservable lost sales. *European Journal of Operational Research* 207(3):1263–1268.
- Igual L, Seguí S (2017) Introduction to data science. *Introduction to Data Science*, 1–4 (Springer).
- Jo J, Bengio Y (2017) Measuring the tendency of cnns to learn surface statistical regularities. *arXiv preprint arXiv:1711.11561* .
- Katz ML, Shapiro C (1994) Systems competition and network effects. *Journal of economic perspectives* 8(2):93–115.
- 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].
- Kos T, Kosar T, Mernik M (2012) Development of data acquisition systems by using a domain-specific modeling language. *Computers in industry* 63(3):181–192.
- 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, Scanail 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).
- 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.
- Schmitz PW (2013) Job design with conflicting tasks reconsidered. *European Economic Review* 57:108–117.
- Schöttner A (2016) Optimal sales force compensation in dynamic settings: Commissions vs. bonuses. *Management Science* 63(5):1529–1544.
- Shy O (2011) A short survey of network economics. *Review of Industrial Organization* 38(2):119–149.
- Spencer CD (1990) *Digital design for computer data acquisition* (Cambridge University Press).
- 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].
- Vincent J (2016) Twitter taught microsoft’s ai chatbot to be a racist asshole in less than a day. <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist> [Online; accessed 11-September-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.

$$\frac{1}{\nu} F^{-1}\left(1 - \frac{\delta}{\nu}\right) - \frac{\delta}{\nu^2} \frac{1}{f\left(F^{-1}\left(1 - \frac{\delta}{\nu}\right)\right)} = 0. \quad (10)$$

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

$$\phi(\xi) = 0. \quad (11)$$

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 Lemma 2. Using Lemma 1, and fixing $\delta_{1\ell}$ and δ_{1h} , we obtain the following problem where we only optimize over the payments $x_{i\ell}$ and x_{ih} for $i = 1, 2$.

$$\min_{\substack{x_{1h}, x_{1\ell} \geq 0 \\ x_{2h}, x_{2\ell} \geq 0}} \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[\pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} + x_{1\ell}] \quad (12)$$

$$\text{s.t.} \quad \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[x_{1\ell} + J_p(d_0)] - \kappa \geq \pi_s(d_0)x_{1h} + \bar{\pi}_s(d_0)[x_{1\ell} + J_p(d_0)] \quad (13)$$

$$\pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} - \kappa \geq \pi_s(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_s(d_0 + \delta_{1\ell})x_{2\ell} \quad (14)$$

$$J_p(d_0) = \pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} - \kappa. \quad (15)$$

To solve this problem, we first consider a subproblem where we optimize over only x_{1h} and $x_{1\ell}$ for a fixed pair of x_{2h} and $x_{2\ell}$ which satisfy (14). Since $\delta_{1\ell}$ is also fixed, we get a constant for $J_p(d_0)$. Thus, the subproblem corresponds to minimizing the expected payments in the first period while satisfying the incentive constraint (13), and given by

$$\min_{x_{1h}, x_{1\ell} \geq 0} \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)x_{1\ell} \quad (16)$$

$$\text{s.t.} \quad \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[x_{1\ell} + J_p(d_0)] - \kappa \geq \pi_s(d_0)x_{1h} + \bar{\pi}_s(d_0)[x_{1\ell} + J_p(d_0)] \quad (17)$$

Following Proposition 4.2 in Laffont and Martimort (2009, p. 157), we get the optimal solution to the subproblem as $\tilde{x}_{1\ell}^* = 0$ and

$$\tilde{x}_{1h}^* = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + J_p(d_0).$$

Note that this optimal solution \tilde{x}_{1h}^* and $\tilde{x}_{1\ell}^*$ is for a fixed value of $x_{2\ell}$ and x_{2h} . Thus, we know optimize over $x_{2\ell}$ and x_{2h} by incorporating the values of $\tilde{x}_{1\ell}^*$ and \tilde{x}_{1h}^* to Problem (12). Arranging terms, we obtain

$$\begin{aligned} \min_{x_{2h}, x_{2\ell} \geq 0} \quad & \pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} + \frac{\kappa}{\rho(d_0)} - \pi_w(d_0)\kappa \\ \text{s.t.} \quad & \pi_w(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_w(d_0 + \delta_{1\ell})x_{2\ell} - \kappa \geq \pi_s(d_0 + \delta_{1\ell})x_{2h} + \bar{\pi}_s(d_0 + \delta_{1\ell})x_{2\ell}. \end{aligned} \quad (18)$$

Invoking the same proposition from Laffont and Martimort (2009) one more time, we get the optimal solution as $x_{2\ell}(\delta_{1\ell}) = 0$ and

$$x_{2h}(\delta_{1\ell}) = \frac{\kappa}{\pi_w(d_0 + \delta_{1\ell}) - \pi_s(d_0 + \delta_{1\ell})}.$$

Evaluating \tilde{x}_{1h}^* using $x_{2h}(\delta_{1\ell})$ and $x_{2\ell}(\delta_{1\ell})$, we find

$$x_{1h}(\delta_{1\ell}) = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + \frac{\kappa}{\rho(d_0 + \delta_{1\ell})} - \kappa \text{ and } x_{1\ell}(\delta_{1\ell}) = 0.$$

The total expected payment follows from using the values of the optimal solution x_{ih} and $x_{i\ell}$ for $i = 1, 2$, and we conclude the proof. Q.E.D.

Proof of Proposition 1. Using Lemma 1 and the optimal payments derived in Lemma 2, we reformulate Problem (6) to a problem where we only maximize over $\delta_{1\ell}$ and δ_{1h} .

$$\begin{aligned} J_{1a}(d_0) = \max_{\delta_{1h}, \delta_{1\ell} \in [0, \nu]} \quad & \pi_w(d_0) \left[R_\nu(\delta_{1h}, \alpha_h) + R_\nu(\delta^M, \alpha_h) - \frac{\kappa}{\rho(d_0 + \delta_{1\ell})} + \kappa \right] - \frac{\kappa}{\rho(d_0)} \\ & + \bar{\pi}_w(d_0) \left[R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})R_\nu(\delta^M, \alpha_h) + \bar{\pi}_w(d_0 + \delta_{1\ell})R_\nu(\delta^M, \alpha_\ell) - \frac{\kappa}{\rho(d_0 + \delta_{1\ell})} \right]. \end{aligned} \quad (19)$$

Note that we can separate the optimization problem above into maximizing over δ_{1h} and $\delta_{1\ell}$. For the part, we maximize over δ_{1h} , we get the myopic problem so the unique optimal solution is given by Lemma 1 as $\delta_{1h} = \delta^M$. Thus, we finally get

$$\begin{aligned} J_{1a}(d_0) = \pi_w(d_0) [2R_\nu(\delta^M, \alpha_h) + \kappa] - \frac{\kappa}{\rho(d_0)} \\ + \max_{\delta_{1\ell} \in [0, \nu]} \bar{\pi}_w(d_0) [R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})R_\nu(\delta^M, \alpha_h) + \bar{\pi}_w(d_0 + \delta_{1\ell})R_\nu(\delta^M, \alpha_\ell)] - \frac{\kappa}{\rho(d_0 + \delta_{1\ell})} \\ = \pi_w(d_0) [2R_\nu(\delta^M, \alpha_h) + \kappa] - \frac{\kappa}{\rho(d_0)} + \bar{\pi}_w(d_0)\tau\alpha_\ell \\ + \max_{\delta_{1\ell} \in [0, \nu]} \bar{\pi}_w(d_0) [R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})\tau(\alpha_h - \alpha_\ell)] - \frac{\kappa}{\rho(d_0 + \delta_{1\ell})}. \end{aligned} \quad (20)$$

The second equality follows from the fact that $R_\nu(\delta^M, \alpha_i) = \tau\alpha_i$ for $i \in \{h, \ell\}$ (see Lemma 1). Here, the uniqueness of $\delta_{1\ell} = \delta^*$ follows because the objective function $\bar{\pi}_w(d_0) [R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})\tau(\alpha_h - \alpha_\ell)] - \frac{\kappa}{\rho(d_0 + \delta_{1\ell})}$ is concave and continuous in $\delta_{1\ell}$ and $[0, \nu]$ is a compact interval. Q.E.D.

Proof of Proposition 2. After removing incentive compatibility constraints and setting all payments to κ in Problem (6), we use Lemma 1 and obtain the following optimization problem

$$J_{1a}^{\text{FB}}(d_0) = \max_{\delta_{1h}, \delta_{1\ell} \in [0, \nu]} \pi_w(d_0)[R_\nu(\delta_{1h}, \alpha_h) + R_\nu(\delta^{\text{M}}, \alpha_h) - \kappa] + \bar{\pi}_w(d_0)[R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})R_\nu(\delta^{\text{M}}, \alpha_h) + \bar{\pi}_w(d_0 + \delta_{1\ell})R_\nu(\delta^{\text{M}}, \alpha_\ell) - 2\kappa] \quad (21)$$

Isolating δ_{1h} and using Lemma 1 again, we get $\delta_{1h} = \delta^{\text{M}}$. The remaining optimization problem is given by

$$J_{1a}^{\text{FB}}(d_0) = \pi_w(d_0)[2R_\nu(\delta^{\text{M}}, \alpha_h) - \kappa] + \bar{\pi}_w(d_0)(\tau\alpha_\ell - 2\kappa) + \bar{\pi}_w(d_0) \left\{ \max_{\delta_{1\ell} \in [0, \nu]} R_\nu(\delta_{1\ell}, \alpha_\ell) + \pi_w(d_0 + \delta_{1\ell})\tau(\alpha_h - \alpha_\ell) \right\}. \quad (22)$$

Because π_w , R_ν are concave and $d_0 + \delta_{1\ell}$ is a linear function of $\delta_{1\ell}$, the objective function inside curly brackets is concave in $\delta_{1\ell}$. This implies that the optimal solution $\delta_{1\ell} = \delta^{\text{FB}}$ is unique.

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

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

Note that the objective function in (22) is concave in $\delta_{1\ell}$, and increasing when evaluated at $\delta_{1\ell} = \delta^{\text{M}}$. Therefore, it follows that $\delta^{\text{FB}} > \delta^{\text{M}}$. Q.E.D.

Proof of Theorem 1. Note that if $\rho(d)$ is constant, then the objective function in Proposition 1 and the one in Proposition 2 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 $\max_{\delta \in [0, \nu]} R_\nu(\delta, \alpha_\ell) + \pi_w(d_0 + \delta)\tau(\alpha_h - \alpha_\ell)$ (see Proposition 2).

$$R_\nu(\delta^{\text{FB}}, \alpha_\ell) + \pi_w(d_0 + \delta^{\text{FB}})\tau(\alpha_h - \alpha_\ell) \geq R_\nu(\delta, \alpha_\ell) + \pi_w(d_0 + \delta)\tau(\alpha_h - \alpha_\ell), \forall \delta \in [0, \nu] \quad (23)$$

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

$$-\frac{\kappa}{\rho(d_0 + \delta^{\text{FB}})} > -\frac{\kappa}{\rho(d_0 + \delta^*)} \quad (24)$$

Inequalities (23) and (24) imply that the objective function $\bar{\pi}_w(d_0)[R_\nu(\delta, \alpha_\ell) + \pi_w(d_0 + \delta)\tau(\alpha_h - \alpha_\ell)] - \frac{\kappa}{\rho(d_0 + \delta)}$ evaluated at $\delta = \delta^{\text{FB}}$ is strictly larger than the value obtained by evaluating the same at $\delta = \delta^*$. Therefore, the condition $\delta^* < \delta^{\text{FB}}$ contradicts with the fact that δ^* is the optimal solution (see Proposition 1). The last item of the theorem can be proved by following the same steps with $\delta^* > \delta^{\text{FB}}$. Q.E.D.

Proof of Proposition 3. In a pricing period (e.g., Period 1b), the total consumer surplus for a given price p and accuracy α is given by

$$\int_0^1 1_{\{\alpha v \geq p\}}(\alpha v - p)f(v)dv = \frac{\alpha}{2} - p + \frac{p^2}{2\alpha},$$

where $1_{\{\cdot\}}$ is the indicator function. Similarly, the total social welfare for a given price p and accuracy α is

$$\int_0^1 1_{\{\alpha v \geq p\}}\alpha v f(v)dv = \frac{\alpha}{2} - \frac{p^2}{2\alpha}.$$

In order to find the total expected social welfare and the customer surplus at equilibrium, we use the equilibrium prices and take expectation with respect to the induced probabilities of the amounts of collected data at equilibrium. The total expected social welfare and the customer surplus are in order.

$$\begin{aligned} W^* &= \pi_w(d_0)\alpha_h \left(1 - (1 - 2\tau)^2\right) + \bar{\pi}_w(d_0)\frac{\alpha_\ell}{2} \left(1 - \left(1 - \frac{\delta^*}{\nu}\right)^2\right) \\ &\quad + \bar{\pi}_w(d_0) \left[\pi_w(d_0 + \delta^*)\frac{\alpha_h}{2} \left(1 - (1 - 2\tau)^2\right) + \bar{\pi}_w(d_0 + \delta^*)\frac{\alpha_\ell}{2} \left(1 - (1 - 2\tau)^2\right) \right] \\ C^* &= \pi_w(d_0)2\tau^2(2\alpha_h - \alpha_\ell) + 2\tau^2\alpha_\ell + \bar{\pi}_w(d_0)\frac{\alpha_\ell(\delta^*)^2}{2\nu^2} + \bar{\pi}_w(d_0)\pi_w(d_0 + \delta^*)2\tau^2(\alpha_h - \alpha_\ell) \end{aligned}$$

The first-best benchmarks W^{FB} and C^{FB} are found by evaluating the terms above using δ^{FB} instead of δ^* . Since π_w is increasing and $\delta^*, \delta^{\text{FB}} \in [0, \nu]$, the monotonicity result of the proposition follows. Q.E.D.

Proof of Proposition 4. From Theorem 1, we know $\delta^* = \delta^{\text{FB}}$ ($\delta^* \geq \delta^{\text{FB}}$ and $\delta^* \leq \delta^{\text{FB}}$) when ρ is constant (increasing and, respectively decreasing). Furthermore, Proposition 3 implies that $W^* = W^{\text{FB}}$ and $C^* = C^{\text{FB}}$ ($W^* \geq W^{\text{FB}}$ and $C^* \geq C^{\text{FB}}$; and $W^* \leq W^{\text{FB}}$ and $C^* \leq C^{\text{FB}}$) when $\delta^* = \delta^{\text{FB}}$ ($\delta^* \geq \delta^{\text{FB}}$ and, respectively $\delta^* \leq \delta^{\text{FB}}$).

Proof of Proposition 6. We use the first-order condition for the maximization problem provided in Proposition 2 and the implicit function theorem to prove this result. The first-order condition is

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

Note here that the left-hand side of the equality above 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. From the implicit function theorem, we know that $\frac{\partial \delta^{\text{FB}}}{\partial \nu}$ has the same sign with

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

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

Proof of Theorem 2. We use the first-order condition for the maximization problem provided in Proposition 1 and the implicit function theorem to prove this result. The first-order condition is

$$\bar{\pi}_w(d_0) \left[\frac{\alpha_\ell}{\nu} \left(1 - \frac{2\delta^*}{\nu} \right) + \pi'_w(d_0 + \delta^*)(\alpha_h - \alpha_\ell)\tau \right] + \frac{\kappa\rho'(d_0 + \delta^*)}{[\rho(d_0 + \delta^*)]^2} = 0. \quad (25)$$

Note 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 $d \in [d_{\min}, d_{\max}]$. 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 2) 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 at an interior ν .

On the other hand, if δ^* crosses $\nu/4$ at some point $v \in (0, d_{\max}/2]$, then δ^* has to be increasing first, i.e., $\delta^* \geq \nu/4$ for all $\nu \in (0, v]$, and then has to be decreasing, i.e., $\delta^* \leq \nu/4$ for all $\nu \in [v, d_{\max}/2]$. This implies $\hat{\nu} = v$. 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.

Proof of Theorem 3. The firm's expected profit derived in Proposition 1 is

$$\begin{aligned} J_{1a}(d_0) = & \pi_w(d_0) [2R_\nu(\delta^{\text{M}}, \alpha_h) + \kappa] - \frac{\kappa}{\rho(d_0)} + \bar{\pi}_w(d_0)\tau\alpha_\ell \\ & + \bar{\pi}_w(d_0) [R_\nu(\delta^*, \alpha_\ell) + \pi_w(d_0 + \delta^*)\tau(\alpha_h - \alpha_\ell)] - \frac{\kappa}{\rho(d_0 + \delta^*)}. \end{aligned}$$

We take derivative of $J_{1a}(d_0)$ with respect to ν . For simplicity, we denote by δ_ν^* the derivative of δ^* with respect to ν . Note that we consider interior optimal solutions δ^* which satisfy the first-order condition so δ_ν^* exists.

$$\begin{aligned} \frac{\partial J_{1a}(d_0)}{\partial \nu} = & \bar{\pi}_w(d_0) \left[\frac{\alpha_\ell}{\nu} \left(\delta_\nu^* - \frac{\delta^*}{\nu} \right) \left(1 - \frac{2\delta^*}{\nu} \right) + \pi_w(d_0 + \delta^*)\tau(\alpha_h - \alpha_\ell)\delta_\nu^* \right] + \frac{\kappa\rho'(d_0 + \delta^*)\delta_\nu^*}{[\rho(d_0 + \delta^*)]^2} \\ = & \delta_\nu^* \left\{ \bar{\pi}_w(d_0) \left[\frac{\alpha_\ell}{\nu} \left(1 - \frac{2\delta^*}{\nu} \right) + \pi_w(d_0 + \delta^*)\tau(\alpha_h - \alpha_\ell) \right] + \frac{\kappa\rho'(d_0 + \delta^*)}{[\rho(d_0 + \delta^*)]^2} \right\} + \frac{\alpha_\ell}{\nu} \left(\frac{\delta^*}{\nu} \right) \left(\frac{2\delta^*}{\nu} - 1 \right) \end{aligned}$$

From the first-order condition for the maximization problem in Proposition 1, and the implicit function theorem, we know that

$$\bar{\pi}_w(d_0) \left[\frac{\alpha_\ell}{\nu} \left(1 - \frac{2\delta^*}{\nu} \right) + \pi'_w(d_0 + \delta^*)(\alpha_h - \alpha_\ell)\tau \right] + \frac{\kappa\rho'(d_0 + \delta^*)}{[\rho(d_0 + \delta^*)]^2} = 0. \quad (26)$$

Thus, it follows that

$$\frac{\partial J_{1a}(d_0)}{\partial \nu} = \frac{\alpha_\ell}{\nu} \left(\frac{\delta^*}{\nu} \right) \left(\frac{2\delta^*}{\nu} - 1 \right) = \frac{\alpha_\ell}{\nu} \left(\frac{\delta^*}{\nu} \right) \frac{2}{\nu} \left(\delta^* - \frac{\nu}{2} \right).$$

The sign of $\partial J_{1a}(d_0)/\partial \nu$ is determined by $\delta^* - \nu/2$. When ρ is increasing $\delta^* \geq \delta^M = \nu/2$, thus $\partial J_{1a}(d_0)/\partial \nu$ is positive.

Recall that $\delta = q\nu$ by definition. Thus, we denote by q^* the quantity corresponding to δ^* , and focus on it to prove the second item of the theorem. We next provide a lemma whose proof is delegated to the end of the proof of this theorem.

LEMMA 3. *When ρ is decreasing, q^* is decreasing in ν if $q^* \leq 1/2$.*

Note also that the sign of $\partial J_{1a}(d_0)/\partial \nu$ is determined by the sign of $(q^* - 1/2)$. When ρ is decreasing δ^* can be larger or smaller than $\delta^M = \nu/2$, i.e., q^* can be larger than smaller or $1/2$. We know that $\partial J_{1a}(d_0)/\partial \nu$ is positive as long as $q^* > 1/2$. If q^* crosses $1/2$ at some $\tilde{\nu}$, then it follows that $\partial J_{1a}(d_0)/\partial \nu$ is negative for all $\nu \geq \tilde{\nu}$ because q^* is decreasing if $q^* \leq 1/2$ when ρ decreasing (see Lemma 3). This concludes the proof. Q.E.D.

Proof of Lemma 3. By the change of variables $\delta = \nu q$, we get the first-order condition for the maximization problem in Proposition 1 in terms of q^* . Recall that we define $q^* = \delta^*/\nu$.

$$\frac{1}{2\alpha_\ell} \pi'_w(d_0 + q^*\nu)(\alpha_h - \alpha_\ell)\tau\nu + \frac{1}{2\bar{\pi}_w(d_0)\alpha_\ell} \frac{\kappa\rho'(d_0 + q^*\nu)\nu}{[\rho(d_0 + q^*\nu)]^2} = q^* - \frac{1}{2}. \quad (27)$$

Using the implicit function theorem, we know that the sign of the derivative of q^* with respect to ν is the same as

$$\underbrace{\left[\bar{\pi}_w(d_0)\pi'_w(d_0 + q^*\nu)\tau(\alpha_h - \alpha_\ell) + \frac{\kappa\rho'(d_0 + q^*\nu)}{[\rho(d_0 + q^*\nu)]^2} \right]}_{(i)} \quad (28)$$

$$+ \underbrace{\left[\bar{\pi}_w(d_0)\pi''_w(d_0 + q^*\nu)\tau(\alpha_h - \alpha_\ell) + \frac{\kappa\rho''(d_0 + q^*\nu)}{[\rho(d_0 + q^*\nu)]^2} - \frac{2\kappa(\rho'(d_0 + q^*\nu))^2}{[\rho(d_0 + q^*\nu)]^3} \right]}_{(ii)} \nu q^*. \quad (29)$$

First note that q^* can take values smaller than $1/2$ when ρ is decreasing. If $q^* \leq 1/2$, the first-order condition (27) implies that (i) is negative. We also know that (ii) is always negative because π_w is concave and $1/\rho$ is convex. Thus, q^* starts to decrease after crossing $1/2$. Q.E.D.

Proof of Proposition 7. Note that the expected profit is decreasing in ν when ρ is decreasing and q^* is smaller than $1/2$. Thus, in this proof we show that q^* is decreasing in d_0 when ρ is decreasing. To do so, we use (27) and the implicit function theorem. The derivative of q^* with respect to d_0 has the same sign with the following term:

$$-\pi'_w(d_0)[\alpha_\ell(1 - 2q^*) + \pi'_w(d_0 + q^*\nu)\tau(\alpha_h - \alpha_\ell)\nu] + \bar{\pi}_w(d_0)\pi''_w(d_0 + q^*\nu)\tau(\alpha_h - \alpha_\ell)\nu + \frac{\kappa\rho''(d_0 + q^*\nu)\nu}{[\rho(d_0 + q^*\nu)]^2} - \frac{2\kappa[\rho'(d_0 + q^*\nu)]\nu}{[\rho(d_0 + q^*\nu)]^3}.$$

Using (27), we obtain

$$\frac{\pi'_w(d_0)\kappa\rho'(d_0+q^*\nu)\nu}{[\rho(d_0+q^*\nu)]^2\bar{\pi}_w(d_0)} + \bar{\pi}_w(d_0)\pi''_w(d_0+q^*\nu)\tau(\alpha_h - \alpha_\ell)\nu + \frac{\kappa\rho''(d_0+q^*\nu)\nu}{[\rho(d_0+q^*\nu)]^2} - \frac{2\kappa[\rho'(d_0+q^*\nu)]\nu}{[\rho(d_0+q^*\nu)]^3}.$$

Here, the expression above is negative because ρ is decreasing, $1/\rho$ is convex; and π_w is an increasing concave function. Q.E.D.

B. Optimal Effort and Price Decisions of the Firm

In this section of the 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 8. *The firm finds it optimal to retrain the algorithm if $\underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) > \kappa/\bar{\pi}_w(d_{\max})$.*

Proof of Proposition 8. Since the firm commits to the contract before the accuracy realizations, we compare the expected profit $J_{1a}(d_0)$ with the expected profit (denoted by $\tilde{J}_{1a}(d_0)$) collected when the firm hires the provider only for the product development. In Proposition 1, we derive $J_{1a}(d_0)$. The expected profit $\tilde{J}_{1a}(d_0)$ for the no-retraining case is given by

$$\tilde{J}_{1a}(d_0) = \pi_w(d_0)2\alpha_h\tau + \bar{\pi}_w(d_0)2\alpha_\ell\tau - \frac{\kappa}{\rho(d_0)}.$$

Because there is no retraining, in both periods the product is sold with the initial accuracy that can be α_h with probability $\pi_w(d_0)$ and α_ℓ with the remaining probability. In each case, it is optimal for the firm to use the myopic price, therefore the expected revenue of the firm is given by $\pi_w(d_0)2\alpha_h\tau + \bar{\pi}_w(d_0)2\alpha_\ell\tau$. The provider is hired only for the first period. In this case, the optimal contract for the provider is such that there is a payment of $\kappa/[\pi_w(d_0) - \pi_s(d_0)]$ in case of a high accuracy, and no payment otherwise. This causes an expected cost of $\kappa/\rho(d_0)$.

The difference between this expected profits is

$$J_{1a}(d_0) - \tilde{J}_{1a}(d_0) = \pi_w(d_0)\kappa + \bar{\pi}_w(d_0) \left[R_\nu(\delta^*, \alpha_\ell) + \pi_w(d_0 + \delta^*)\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\bar{\pi}(d_0)\rho(d_0 + \delta^*)} - \tau\alpha_\ell \right].$$

In order to find a sufficient condition for $J_{1a}(d_0) - \tilde{J}_{1a}(d_0)$ to be positive, one can focus on the term inside the square brackets above because the remaining terms are all positive. Thus, we next show that the second term in the square brackets is larger than $\alpha_\ell\tau$ when $\underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) > \kappa/\bar{\pi}_w(d_{\max})$ by finding a uniform positive lower bound for $\pi_w(d)\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\bar{\pi}_w(d)\rho(d)}$ for all $d \in [d_{\min}, d_{\max}]$. Let $\epsilon \triangleq \underline{\Delta\pi}\tau(\alpha_h - \alpha_\ell) - \kappa/\bar{\pi}_w(d_{\max})$, thus $\epsilon > 0$.

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

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

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

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

Proof of Proposition 9. This result follows from the fact that the derivatives of the objective functions in the maximization problems in Proposition 2 and respectively Proposition 1 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 2 and 1.)

Note that Proposition 2 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 Proposition 2 as follows:

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

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 Proposition 1 with respect to δ_1 is given as follows:

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

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 9 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.

In case of the corner solution such that the δ^* takes the value of its upper bound ν , the corresponding price of the product p^* becomes 0 per Section 2. In this case, the firm distributes the product for free to collect the largest available data. We conclude this section of the appendix by providing a necessary and sufficient condition for $\delta^* = \nu$.

LEMMA 4. *The optimal data collection size δ^* equals to ν if and only if the following condition holds.*

$$\bar{\pi}_w(d_0)\tau(\alpha_h - \alpha_\ell)\pi'_w(d_0 + \nu) + \frac{\kappa\rho'(d_0 + \nu)}{[\rho(d_0 + \nu)]^2} \geq \frac{\bar{\pi}_w(d_0)\alpha_\ell}{\nu} \quad (32)$$

Proof of Lemma 4. The firm's optimal data collection size is found by maximizing the objective function in Proposition 1. When the inequality in (32) holds, the objective function is nondecreasing for any $\delta \in [0, \nu]$ because it is concave. Thus, the firm selects the largest data collection size in this case.

C. Contract Termination after Low Accuracy

In this section of the appendix, we consider the setting where the firm can terminate the relation with the provider after observing a low accuracy in the first period, and contract with another provider in the second period. Recall that the firm does not need to retrain the algorithm after a high accuracy in the first period. Thus, firing after a low accuracy terminates the contract with the provider in the first period in any case. The firm then solves two separate contracting problems between which it collects data. To distinguish from the setting with a single provider, in this appendix we denote by θ data collection. The next proposition based on Laffont and Martimort (2009) characterizes the optimal payments of the firm.

PROPOSITION 10. *In case of contracting with two providers in different periods, the optimal payments of the firm are such that $x_{1\ell}^* = 0$, $x_{2\ell}^* = 0$; and*

$$x_{1h}^* = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} \text{ and } x_{2h}^* = \frac{\kappa}{\pi_w(d_0 + \theta^*) - \pi_s(d_0 + \theta^*)}, \quad (33)$$

where θ^* is the unique optimal solution of $\max_{\theta \in [0, \nu]} R_\nu(\theta, \alpha_\ell) - \pi_w(d_0 + \theta)\tau(\alpha_h - \alpha_\ell) - \kappa/\rho(d_0 + \theta)$.

Proof of Proposition 10. The optimal payments of the contracting problem with the first provider is characterized in Proposition 4.2 in Laffont and Martimort (2009). For any data collection θ after a low accuracy in the first period, the firm faces the same contracting problem in the second period but this time with a different accuracy probabilities: $\pi_w(d_0 + \theta)$ and $\pi_s(d_0 + \theta)$. Therefore, the same proposition in Laffont and Martimort (2009) characterizes the optimal payments in the second period for any given θ .

At the end of the first period, the firm maximizes the next periods profit by choosing the optimal data collection size by anticipating the optimal payments when contracting with a new provider in case of obtaining low accuracy in the first period. Otherwise, goes with the myopic price for high accuracy. The firm's optimization problem after the low accuracy is as follows.

$$\max_{\theta \in [0, \nu]} \underbrace{R_\nu(\theta, \alpha_\ell)}_{\text{Current Revenue}} + \underbrace{\pi_w(d_0 + \theta)\tau(\alpha_h - \alpha_\ell)}_{\text{Continuation Revenue}} - \underbrace{\kappa/\rho(d_0 + \theta)}_{\text{Expected Agency Cost}}. \quad (34)$$

Because all terms in the above optimization problem are concave functions of θ and $[0, \nu]$ is a compact set, there exists a unique θ^* . Q.E.D.

As in the case of the single provider we restrict attention to $\theta^* \in (0, \nu)$ for the following results. Note that the objective function in the optimization problem in (34) differs from the one in Proposition 1 by a scaling term $\bar{\pi}_w(d_0)$ for the first two terms: current period revenue and continuation revenue. Therefore, Theorems 1 and 2 can be repeated verbatim using (34). Thus, Proposition 4 is extended to this setting as well. The firm's expected profit in this setting is differently from Theorem 3 given by

$$\begin{aligned} & \pi_w(d_0) [2R_\nu(\delta^M, \alpha_h) + \kappa] - \frac{\kappa}{\rho(d_0)} + \bar{\pi}_w(d_0)\tau\alpha_\ell \\ & + \bar{\pi}_w(d_0) \left[R_\nu(\theta^*, \alpha_\ell) + \pi_w(d_0 + \theta^*)\tau(\alpha_h - \alpha_\ell) - \frac{\kappa}{\rho(d_0 + \theta^*)} \right]. \end{aligned} \quad (35)$$

Here, the difference is that the firm incurs the expected agency cost in the second period only if the first period yields low accuracy. Nevertheless the same steps in the proof of Theorem 3 (including Lemma 3) can be followed verbatim because θ^* is the interior optimal solution to (34) as δ^* is the interior optimal solution to the maximization problem in Proposition 1. To show Proposition 7 is also extended to this setting, we need to prove that q^* is decreasing in d_0 when ρ is decreasing in this setting, too.

PROPOSITION 11. *When ρ is decreasing, q^* is decreasing in d_0 .*

Proof of Proposition 11. As in the proof of Proposition 7, we use the first-order condition for (34) and the implicit function theorem. The derivative of q^* with respect to d_0 has the same sign with the following term:

$$\pi_w''(d_0 + \theta^*)\nu(\alpha_h - \alpha_\ell)\tau + \frac{\kappa\rho''(d_0 + \theta^*)\nu}{[\rho(d_0 + \theta^*)]^2} - \frac{2\kappa[\rho'(d_0 + \theta^*)]^2\nu}{[\rho(d_0 + \theta^*)]^3} \quad (36)$$

Because π_w is concave and $1/\rho$ is convex it follows that the term above is negative. Q.E.D.

We conclude this appendix by providing the proof of Proposition 5.

Proof of Proposition 5. This proposition stems from the fact that θ^* is larger than δ^* when ρ is decreasing and vice versa. Specifically, we show in Proposition 3 that collecting more data increases both the social welfare and the customer surplus. Thus, to compare those societal benefits we can directly compare the optimal data collection sizes in two settings. The optimal data collection δ^* and θ^* satisfy the following first-order conditions.

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

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

If ρ is increasing, then $\alpha_\ell \left(\frac{1}{\nu} - \frac{2\theta^*}{\nu^2} \right) + \tau(\alpha_h - \alpha_\ell) \pi'_w(d_0 + \theta^*) + \frac{\kappa \rho'(d_0 + \theta^*)}{\bar{\pi}_w(d_0) [\rho(d_0 + \theta^*)]^2} > 0$ because $\bar{\pi}_w(d_0) < 1$. Hence, $\theta^* < \delta^*$ because both objective functions are concave, hence their derivatives are decreasing. This implies the firm collects more data if the provider cannot be fired after a low accuracy in the first period and the data impact is increasing. In the case of decreasing ρ , $\alpha_\ell \left(\frac{1}{\nu} - \frac{2\theta^*}{\nu^2} \right) + \tau(\alpha_h - \alpha_\ell) \pi'_w(d_0 + \theta^*) + \frac{\kappa \rho'(d_0 + \theta^*)}{\bar{\pi}_w(d_0) [\rho(d_0 + \theta^*)]^2} < 0$ so $\theta^* > \delta^*$. Q.E.D.

D. No Commitment across Periods

In this section of the appendix, we consider a setting where the firm has no commitment power across periods when contracting with the provider. In order to differentiate the data collection sizes in this section from δ in our paper, we denote by γ data collection sizes. Furthermore, we use the hat symbol $\hat{\cdot}$ over the value functions for this setting.

Since the firm has no commitment power, we formulate the firm's problem using backward induction starting from the last period. Denote then by $\hat{J}_{2b}(\alpha_2, d_1)$ the firm's optimal expected profit in Period 2b given accuracy α_2 and data size d_1 , such that

$$\hat{J}_{2b}(\alpha_2, d_1) = \max_{\delta_2 \in [0, \nu]} R_\nu(\delta_2, \alpha_2) \quad (37)$$

The firm chooses the amount of collected data (or equivalently the price) so as to maximize the expected revenue in the current period as in (2).

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

$$\hat{J}_{2a}(\alpha_\ell, d_1) = \max_{x_{2h}, x_{2\ell} \geq 0} \pi_w(d_1) [\hat{J}_{2b}(\alpha_h, d_1) - x_{2h}] + \bar{\pi}_w(d_1) [\hat{J}_{2b}(\alpha_\ell, d_1) - x_{2\ell}] \quad (38)$$

s.t.

$$\pi_w(d_1) x_{2h} + \bar{\pi}_w(d_1) x_{2\ell} - \kappa \geq \pi_s(d_1) x_{2h} + \bar{\pi}_s(d_1) x_{2\ell} \quad (39)$$

$$\hat{J}_{2a}(\alpha_h, d_1) = \hat{J}_{2b}(\alpha_h, d_1). \quad (40)$$

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 (39). These payments are then deduced 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 (38).

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

$$\hat{J}_{1b}(\alpha_1, d_0) = \max_{\gamma_1 \in [0, \nu]} R_\nu(\gamma_1, \alpha_1) + \hat{J}_{2a}(\alpha_1, d_0 + \gamma_1) \text{ for } \alpha_1 \in \{\alpha_h, \alpha_\ell\}. \quad (41)$$

In contrast to Problem (37), 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 γ_1 (or equivalently price) affects current revenues directly and future ones indirectly by increasing the dataset size to $d_0 + \gamma_1$. We refer to γ^* as the optimal solution of Problem (41) for $\alpha_1 = \alpha_\ell$. When $\alpha_1 = \alpha_h$, no retraining is required and the optimal price is equal to the myopic price.

Given initial data size d_0 , we denote by $\hat{J}_{1a}(d_0)$ the optimal total expected profit in Period 1a, such that

$$\hat{J}_{1a}(d_0) = \max_{x_{1h}, x_{1\ell} \geq 0} \pi_w(d_0)[\hat{J}_{1b}(\alpha_h, d_0) - x_{1h}] + \bar{\pi}_w(d_0)[\hat{J}_{1b}(\alpha_\ell, d_0) - x_{1\ell}] \quad (42)$$

s.t.

$$\pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)[x_{1\ell} + \hat{J}_p(d_0)] - \kappa \geq \pi_s(d_0)x_{1h} + \bar{\pi}_s(d_0)[x_{1\ell} + \hat{J}_p(d_0)] \quad (43)$$

$$\hat{J}_p(d_0) = \pi_w(d_0 + \delta^*)x_{2h}^* + \bar{\pi}_w(d_0 + \delta^*)x_{2\ell}^* - \kappa. \quad (44)$$

The optimization problem of \hat{J}_{2a} in (38) 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 12. *Given dataset size d_1 , the unique optimal payments for Problem (38) 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*

$$\hat{J}_{2a}(\alpha_\ell, d_1) = \pi_w(d_1)\tau\alpha_h + \bar{\pi}_w(d_1)\tau\alpha_\ell - \frac{\kappa}{\rho(d_1)} \text{ and } \hat{J}_{2a}(\alpha_h, d_1) = \tau\alpha_h. \quad (45)$$

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

Proof of Proposition 12. Following Proposition 4.2 in Laffont and Martimort (2009, p. 157), we obtain the optimal payments for Problem (38) $x_{2\ell}^* = 0$ 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

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

If $\alpha_1 = \alpha_h$, there are no payments so $\hat{J}_{2a}(\alpha_h, d_1) = \hat{J}_{2b}(\alpha_h, d_1) = \tau\alpha_h$.

Q.E.D.

The optimal expected continuation profit $\hat{J}_{1b}(\alpha_1, d_0)$ in Problem (41) is then obtained by using expected profit $\hat{J}_{2a}(\alpha_1, d_0 + \gamma)$ from (45) in Proposition 12. We show next that the optimal data size solving Problem (41) and hence the corresponding optimal price are unique

PROPOSITION 13. *The optimal solution to Problem (41) is unique and equal to δ^M if $\alpha_1 = \alpha_h$.*

Proof of Proposition 13. Using the expected profit \hat{J}_{2a} derived in (45), we write Problem (41) as follows.

$$\max_{\gamma_1 \in [0, \nu]} R_\nu(\gamma_1, \alpha_1) + 1\{\alpha_1 = \alpha_\ell\} \left[\pi_w(d_0 + \gamma_1)\alpha_h\tau + \bar{\pi}_w(d_0 + \gamma_1)\alpha_\ell\tau - \frac{\kappa}{\rho(d_0 + \gamma_1)} \right] \quad (46)$$

If $\alpha_1 = \alpha_h$, it is straightforward to see from (46) that the unique optimal solution to Problem (41) is δ^M . Otherwise, the optimal solution γ^* solves the following problem

$$\max_{\gamma_1 \in [0, \nu]} R_\nu(\gamma_1, \alpha_\ell) + \pi_w(d_0 + \gamma_1)(\alpha_h - \alpha_\ell)\tau - \frac{\kappa}{\rho(d_0 + \gamma_1)} + \alpha_\ell\tau \quad (47)$$

Because π_w , R_ν and $-1/\rho$ are concave, and $d_0 + \gamma_1$ is a linear function of γ_1 , the objective function in (47) is concave in γ_1 . This implies that the optimal solution γ^* is unique. Q.E.D.

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

PROPOSITION 14. *Given initial data size d_0 , the unique optimal payments for Problem (6) are*

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

Proof of Proposition 14. In Period 1a, the optimization problem of the firm is

$$\begin{aligned} \hat{J}_{1a}(d_0) = \pi_w(d_0)\hat{J}_{1b}(\alpha_h, d_0) + \bar{\pi}_w(d_0)\hat{J}_{1b}(\alpha_\ell, d_0) - \min_{x_{1h}, x_{1\ell} \geq 0} \pi_w(d_0)x_{1h} + \bar{\pi}_w(d_0)x_{1\ell} \\ \text{st. (43), (44)} \end{aligned}$$

Using the optimal payments x_{2h}^* and $x_{2\ell}^*$ derived in Proposition 12, we evaluate the expected continuation profit $\hat{J}_p(d_0)$ of the provider, and $\hat{J}_p(d_0) = \frac{\kappa}{\rho(d_0 + \gamma^*)} - \kappa$. Characterizing $\hat{J}_p(d_0)$, we reduce this problem to a standard principal-agent model where the cost of effort is $\kappa + [\pi_w(d_0) - \pi_s(d_0)]\hat{J}_p(d_0)$. Therefore, Proposition 4.2 in Laffont and Martimort (2009, p. 157) implies that the optimal solution of the optimization problem in $\hat{J}_{1b}(d_0)$ is $x_{1\ell}^* = 0$ and $x_{1h}^* = \frac{\kappa}{\pi_w(d_0) - \pi_s(d_0)} + \frac{\kappa}{\rho(d_0 + \gamma^*)} - \kappa$. Q.E.D.

When focusing on $\gamma^* \in (0, \nu)$, we have that the objective function in the optimization problem in (47) differs from the one in Proposition 1 by a scaling term $\bar{\pi}_w(d_0)$ for the first two terms: current period revenue and continuation revenue. Therefore, Theorems 1 and 2 can be repeated verbatim using (47). Thus, Proposition 4 is extended to this setting as well.

Recent ESMT Working Papers

	ESMT No.
Queueing systems with rationally inattentive customers Caner Canyakmaz, ESMT Berlin Tamer Boyaci, ESMT Berlin	18-04 (R1)
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 Hans W. Friederiszick, ESMT Berlin and E.CA Economics Linda Gratz, E.CA Economics Michael Rauber, E.CA Economics	19-03 (R1)
Beyond retail stores: Managing product proliferation along the supply chain Işık Biçer, Schulich School of Business, York University Florian Lückner, Cass Business School, City, University of London Tamer Boyaci, ESMT Berlin	19-02 (R1)
Marginality, dividends, and the value in games with externalities Frank Huettner, ESMT Berlin André Casajus, HHL Leipzig Graduate School of Management	19-01
Consumer choice under limited attention when alternatives have different information costs Frank Huettner, ESMT Berlin Tamer Boyaci, ESMT Berlin Yalçın Akçay, Melbourne Business School	16-04 (R3)